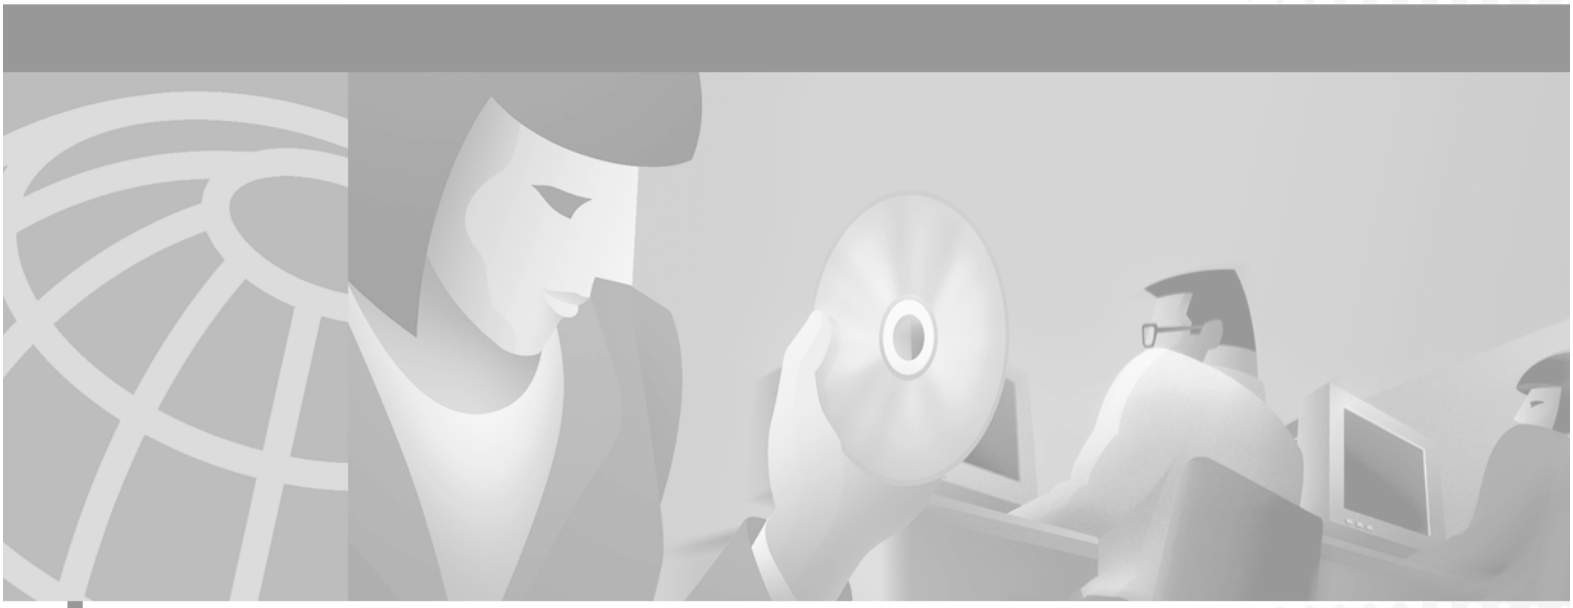

Catalyst 3550 Multilayer Switch Software Configuration Guide

Cisco IOS Release 12.1(8)EA1
February 2002

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: DOC-7811194=
Text Part Number: 78-11194-03

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AccessPath, AtmDirector, Browse with Me, CCIP, CCSI, CD-PAC, *CiscoLink*, the Cisco *Powered* Network logo, Cisco Systems Networking Academy, the Cisco Systems Networking Academy logo, Cisco Unity, Fast Step, Follow Me Browsing, FormShare, FrameShare, IGX, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ Logo, iQ Net Readiness Scorecard, MGX, the Networkers logo, ScriptBuilder, ScriptShare, SMARTnet, TransPath, Voice LAN, Wavelength Router, and WebViewer are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and Discover All That's Possible are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, GigaStack, IOS, IP/TV, LightStream, MICA, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0110R)

Catalyst 3550 Multilayer Switch Software Configuration Guide
Copyright © 2002, Cisco Systems, Inc.
All rights reserved.

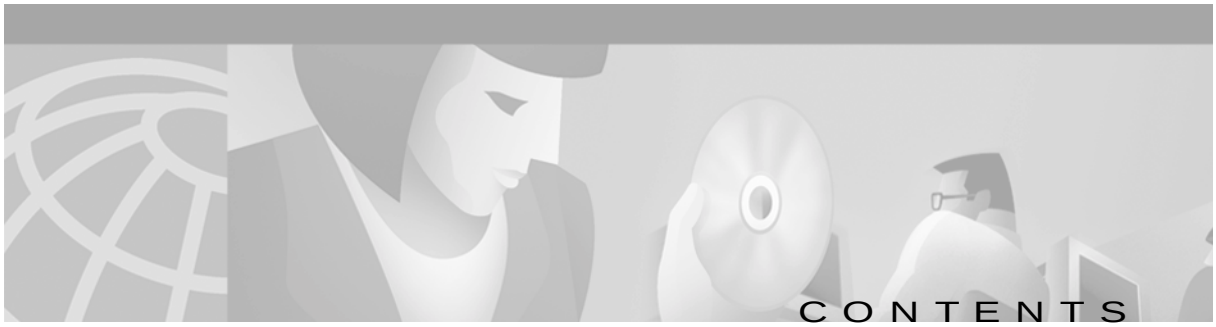

CONTENTS

Preface xxvii

- Audience xxvii
- Purpose xxvii
- Organization xxviii
- Conventions xxx
- Related Publications xxxi
- Obtaining Documentation xxxi
  - World Wide Web xxxi
  - Documentation CD-ROM xxxi
  - Ordering Documentation xxxii
    - Documentation Feedback xxxii
- Obtaining Technical Assistance xxxii
  - Cisco.com xxxiii
  - Technical Assistance Center xxxiii
    - Cisco TAC Web Site xxxiii
    - Cisco TAC Escalation Center xxxiv

CHAPTER 1 **Overview** 1-1

- Features 1-1
- Management Options 1-5
  - Management Interface Options 1-5
  - Advantages of Using CMS and Clustering Switches 1-6
- Network Configuration Examples 1-7
  - Design Concepts 1-7
  - Small to Medium-Sized Network Using Mixed Switches 1-11
  - Large Network Using Only Catalyst 3550 Switches 1-13
  - Multidwelling Network Using Catalyst 3550 Switches 1-14

CHAPTER 2 **Using the Command-Line Interface** 2-1

- IOS Command Modes 2-1
- Getting Help 2-3
- Abbreviating Commands 2-3
- Using no and default Forms of Commands 2-4

Understanding CLI Messages 2-4
Using Command History 2-5
Changing the Command History Buffer Size 2-5
Recalling Commands 2-5
Disabling the Command History Feature 2-5
Using Editing Features 2-6
Enabling and Disabling Editing Features 2-6
Editing Commands through Keystrokes 2-6
Editing Command Lines that Wrap 2-8
Searching and Filtering Output of show and more Commands 2-8
Accessing the CLI 2-9

CHAPTER 3 **Getting Started with CMS 3-1**
Features 3-2
Front Panel View 3-4
Cluster Tree 3-5
Front-Panel Images 3-6
Redundant Power System LED 3-7
Port Modes and LEDs 3-8
VLAN Membership Modes 3-9
Topology View 3-10
Topology Icons 3-12
Device and Link Labels 3-13
Colors in the Topology View 3-14
Topology Display Options 3-14
Menus and Toolbar 3-15
Menu Bar 3-15
Toolbar 3-21
Front Panel View Popup Menus 3-22
Device Popup Menu 3-22
Port Popup Menu 3-22
Topology View Popup Menus 3-23
Link Popup Menu 3-23
Device Popup Menus 3-24
Interaction Modes 3-26
Guide Mode 3-26
Expert Mode 3-26
Wizards 3-26

Tool Tips 3-27
Online Help 3-27
CMS Window Components 3-28
Host Name List 3-28
Tabs, Lists, and Tables 3-29
Icons Used in Windows 3-29
Buttons 3-29
Accessing CMS 3-30
Access Modes in CMS 3-31
HTTP Access to CMS 3-31
Verifying Your Changes 3-32
Change Notification 3-32
Error Checking 3-32
Saving Your Changes 3-32
Using Different Versions of CMS 3-33
Where to Go Next 3-33

CHAPTER 4 **Assigning the Switch IP Address and Default Gateway 4-1**
Understanding the Boot Process 4-1
Assigning Switch Information 4-2
Default Switch Information 4-3
Understanding DHCP-Based Autoconfiguration 4-3
DHCP Client Request Process 4-4
Configuring the DHCP Server 4-5
Configuring the TFTP Server 4-5
Configuring the DNS 4-6
Configuring the Relay Device 4-6
Obtaining Configuration Files 4-7
Example Configuration 4-8
Manually Assigning IP Information 4-10
Checking and Saving the Running Configuration 4-10
Modifying the Startup Configuration 4-12
Default Boot Configuration 4-12
Automatically Downloading a Configuration File 4-12
Specifying the Filename to Read and Write the System Configuration 4-13
Booting Manually 4-13
Booting a Specific Software Image 4-14
Controlling Environment Variables 4-15

Scheduling a Reload of the Software Image 4-17
Configuring a Scheduled Reload 4-17
Displaying Scheduled Reload Information 4-18

CHAPTER 5 **Clustering Switches 5-1**

Understanding Switch Clusters 5-2
Command Switch Characteristics 5-2
Standby Command Switch Characteristics 5-3
Candidate and Member Switches Characteristics 5-3
Planning a Switch Cluster 5-4
Automatic Discovery of Cluster Candidates and Members 5-4
Discovery through CDP Hops 5-5
Discovery through Non-CDP-Capable and Noncluster-Capable Devices 5-6
Discovery through Different VLANs 5-7
Discovery through the Same Management VLAN 5-8
Discovery through Different Management VLANs 5-9
Discovery through Routed Ports 5-10
Discovery of Newly Installed Switches 5-11
HSRP and Standby Command Switches 5-12
Virtual IP Addresses 5-13
Automatic Recovery of Cluster Configuration 5-13
Considerations for Cluster Standby Groups 5-14
IP Addresses 5-15
Host Names 5-16
Passwords 5-16
SNMP Community Strings 5-16
TACACS+ 5-17
Access Modes in CMS 5-17
LRE Profiles 5-17
Availability of Switch-Specific Features in Switch Clusters 5-18
Creating a Switch Cluster 5-18
Enabling a Command Switch 5-19
Adding Member Switches 5-20
Creating a Cluster Standby Group 5-22
Verifying a Switch Cluster 5-24
Using the CLI to Manage Switch Clusters 5-25
Catalyst 1900 and Catalyst 2820 CLI Considerations 5-25
Using SNMP to Manage Switch Clusters 5-26

CHAPTER 6 **Administering the Switch** 6-1

Preventing Unauthorized Access to Your Switch 6-1

Protecting Access to Privileged EXEC Commands 6-2

Default Password and Privilege Level Configuration 6-3

Setting or Changing a Static Enable Password 6-3

Protecting Enable and Enable Secret Passwords with Encryption 6-4

Disabling Password Recovery 6-5

Setting a Telnet Password for a Terminal Line 6-6

Configuring Username and Password Pairs 6-7

Configuring Multiple Privilege Levels 6-8

Setting the Privilege Level for a Command 6-8

Changing the Default Privilege Level for Lines 6-9

Logging into and Exiting a Privilege Level 6-10

Controlling Switch Access with TACACS+ 6-10

Understanding TACACS+ 6-10

TACACS+ Operation 6-12

Configuring TACACS+ 6-13

Default TACACS+ Configuration 6-13

Identifying the TACACS+ Server Host and Setting the Authentication Key 6-13

Configuring TACACS+ Login Authentication 6-14

Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services 6-16

Starting TACACS+ Accounting 6-17

Displaying the TACACS+ Configuration 6-17

Controlling Switch Access with RADIUS 6-17

Understanding RADIUS 6-18

RADIUS Operation 6-19

Configuring RADIUS 6-19

Default RADIUS Configuration 6-20

Identifying the RADIUS Server Host 6-20

Configuring RADIUS Login Authentication 6-23

Defining AAA Server Groups 6-24

Configuring RADIUS Authorization for User Privileged Access and Network Services 6-26

Starting RADIUS Accounting 6-27

Configuring Settings for All RADIUS Servers 6-28

Configuring the Switch to Use Vendor-Specific RADIUS Attributes 6-28

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication 6-29

Displaying the RADIUS Configuration 6-30

Configuring the Switch for Local Authentication and Authorization 6-31

Managing the System Time and Date 6-32
Understanding the System Clock 6-32
Understanding Network Time Protocol 6-32
Configuring NTP 6-34
Default NTP Configuration 6-35
Configuring NTP Authentication 6-35
Configuring NTP Associations 6-36
Configuring NTP Broadcast Service 6-37
Configuring NTP Access Restrictions 6-38
Configuring the Source IP Address for NTP Packets 6-40
Displaying the NTP Configuration 6-41
Configuring Time and Date Manually 6-41
Setting the System Clock 6-42
Displaying the Time and Date Configuration 6-42
Configuring the Time Zone 6-43
Configuring Summer Time (Daylight Saving Time) 6-44
Configuring a System Name and Prompt 6-46
Default System Name and Prompt Configuration 6-46
Configuring a System Name 6-46
Configuring a System Prompt 6-47
Understanding DNS 6-47
Default DNS Configuration 6-48
Setting Up DNS 6-48
Displaying the DNS Configuration 6-49
Creating a Banner 6-49
Default Banner Configuration 6-49
Configuring a Message-of-the-Day Login Banner 6-50
Configuring a Login Banner 6-51
Managing the MAC Address Table 6-51
Building the Address Table 6-52
MAC Addresses and VLANs 6-52
Default MAC Address Table Configuration 6-53
Changing the Address Aging Time 6-53
Removing Dynamic Address Entries 6-54
Configuring MAC Address Notification Traps 6-54
Adding and Removing Static Address Entries 6-56
Displaying Address Table Entries 6-57
Optimizing System Resources for User-Selected Features 6-57
Using the Templates 6-59

CHAPTER 7 **Configuring 802.1X Port-Based Authentication** 7-1

Understanding 802.1X Port-Based Authentication 7-1
Device Roles 7-2
Authentication Initiation and Message Exchange 7-3
Ports in Authorized and Unauthorized States 7-4
Supported Topologies 7-4
Configuring 802.1X Authentication 7-5
Default 802.1X Configuration 7-6
802.1X Configuration Guidelines 7-7
Enabling 802.1X Authentication 7-8
Configuring the Switch-to-RADIUS-Server Communication 7-9
Enabling Periodic Re-Authentication 7-10
Manually Re-Authenticating a Client Connected to a Port 7-11
Changing the Quiet Period 7-11
Changing the Switch-to-Client Retransmission Time 7-12
Setting the Switch-to-Client Frame-Retransmission Number 7-13
Enabling Multiple Hosts 7-13
Resetting the 802.1X Configuration to the Default Values 7-14
Displaying 802.1X Statistics and Status 7-14

CHAPTER 8 **Configuring Interface Characteristics** 8-1

Understanding Interface Types 8-1
Port-Based VLANs 8-2
Switch Ports 8-2
Access Ports 8-2
Trunk Ports 8-3
EtherChannel Port Groups 8-3
Switch Virtual Interfaces 8-4
Routed Ports 8-4
Connecting Interfaces 8-5
Using the Interface Command 8-6
Procedures for Configuring Interfaces 8-7
Configuring a Range of Interfaces 8-9
Configuring and Using Interface Range Macros 8-11
Configuring Layer 2 Interfaces 8-12
Default Layer 2 Ethernet Interface Configuration 8-13
Configuring Interface Speed and Duplex Mode 8-14
Configuration Guidelines 8-14
Setting the Interface Speed and Duplex Parameters 8-14

Configuring IEEE 802.3X Flow Control 8-16
Adding a Description for an Interface 8-17
Monitoring and Maintaining the Layer 2 Interface 8-18
Monitoring Interface and Controller Status 8-18
Clearing and Resetting Interfaces and Counters 8-20
Shutting Down and Restarting the Interface 8-21
Configuring Layer 3 Interfaces 8-22

CHAPTER 9 **Creating and Maintaining VLANs 9-1**
Understanding VLANs 9-1
Number of Supported VLANs 9-2
VLAN Port Membership Modes 9-3
Using the VLAN Trunking Protocol 9-3
The VTP Domain and VTP Modes 9-4
VTP Advertisements 9-5
VTP Version 2 9-6
VTP Pruning 9-6
Configuring VTP 9-8
Default VTP Configuration 9-8
VTP Configuration Guidelines 9-8
Configuring a VTP Server 9-10
Configuring a VTP Client 9-11
Disabling VTP (VTP Transparent Mode) 9-11
Enabling VTP Version 2 9-12
Enabling VTP Pruning 9-13
Monitoring VTP 9-13
VLANs in the VTP Database 9-15
Token Ring VLANs 9-15
Default VLAN Configuration 9-15
VLAN Configuration Guidelines 9-16
Configuring VLANs in the VTP Database 9-17
Adding an Ethernet VLAN 9-17
Modifying an Ethernet VLAN 9-18
Deleting a VLAN from the Database 9-18
Assigning Static-Access Ports to a VLAN 9-19
Displaying VLANs in the VTP Database 9-21
Understanding VLAN Trunks 9-22
Trunking Overview 9-22
Encapsulation Types 9-23

802.1Q Configuration Considerations 9-24
Default Layer 2 Ethernet Interface VLAN Configuration 9-24
Configuring an Ethernet Interface as a Trunk Port 9-25
Configuring a Trunk Port 9-25
Defining the Allowed VLANs on a Trunk 9-27
Changing the Pruning-Eligible List 9-28
Configuring the Native VLAN for Untagged Traffic 9-29
Load Sharing Using STP 9-29
Load Sharing Using STP Port Priorities 9-30
Configuring STP Port Priorities and Load Sharing 9-30
Load Sharing Using STP Path Cost 9-32
Configuring STP Path Costs and Load Sharing 9-32
Understanding VMPS 9-33
Dynamic Port VLAN Membership 9-34
VMPS Database Configuration File 9-34
VMPS Configuration Guidelines 9-36
Default VMPS Configuration 9-37
Configuring an Interface as a Layer 2 Dynamic Access Port 9-37
Entering the IP Address of the VMPS 9-37
Configuring Dynamic Access Ports on VMPS Clients 9-38
Reconfirming VLAN Memberships 9-39
Changing the Reconfirmation Interval 9-39
Changing the Retry Count 9-39
Administering and Monitoring the VMPS 9-40
Troubleshooting Dynamic Port VLAN Membership 9-40
Dynamic Port VLAN Membership Configuration Example 9-40

CHAPTER 10 **Configuring STP** 10-1
Understanding Basic STP Features 10-1
Supported STP Instances 10-2
STP Overview 10-2
Bridge ID, Switch Priority, and Extended System ID 10-3
Election of the Root Switch 10-3
Bridge Protocol Data Units 10-4
STP Timers 10-5
Creating the STP Topology 10-5
STP Interface States 10-6
Blocking State 10-7
Listening State 10-7

Learning State 10-7
Forwarding State 10-8
Disabled State 10-8
STP Address Management 10-8
STP and IEEE 802.1Q Trunks 10-8
VLAN-Bridge STP 10-9
STP and Redundant Connectivity 10-9
Accelerated Aging to Retain Connectivity 10-10
Understanding Advanced STP Features 10-10
Understanding Port Fast 10-10
Understanding BPDU Guard 10-11
Understanding UplinkFast 10-12
Understanding Cross-Stack UplinkFast 10-13
How CSUF Works 10-14
Events that Cause Fast Convergence 10-15
Limitations 10-16
Connecting the Stack Ports 10-16
Understanding BackboneFast 10-18
Understanding Root Guard 10-20
Understanding EtherChannel Guard 10-20
Configuring Basic STP Features 10-21
Default STP Configuration 10-21
Disabling STP 10-22
Configuring the Root Switch 10-22
Configuring a Secondary Root Switch 10-24
Configuring STP Port Priority 10-26
Configuring STP Path Cost 10-27
Configuring the Switch Priority of a VLAN 10-28
Configuring the Hello Time 10-29
Configuring the Forwarding-Delay Time for a VLAN 10-29
Configuring the Maximum-Aging Time for a VLAN 10-30
Configuring STP for Use in a Cascaded Stack 10-30
Displaying STP Status 10-31
Configuring Advanced STP Features 10-32
Configuring Port Fast 10-32
Configuring BPDU Guard 10-33
Configuring UplinkFast for Use with Redundant Links 10-34
Configuring Cross-Stack UplinkFast 10-35
Configuring BackboneFast 10-36

Configuring Root Guard 10-36
Enabling EtherChannel Guard 10-37

CHAPTER 11 **Configuring IGMP Snooping and MVR 11-1**

Understanding IGMP Snooping 11-1
Joining a Multicast Group 11-2
Leaving a Multicast Group 11-4
Immediate-Leave Processing 11-4
Configuring IGMP Snooping 11-5
Default IGMP Snooping Configuration 11-5
Enabling or Disabling IGMP Snooping 11-5
Setting the Snooping Method 11-6
Configuring a Multicast Router Port 11-7
Configuring a Host Statically to Join a Group 11-8
Enabling IGMP Immediate-Leave Processing 11-9
Displaying IGMP Snooping Information 11-9
Understanding Multicast VLAN Registration 11-12
Using MVR in a Multicast Television Application 11-12
Configuring MVR 11-14
Configuration Guidelines and Limitations 11-14
Default MVR Configuration 11-15
Configuring MVR Global Parameters 11-15
Configuring MVR Interfaces 11-16
Displaying MVR Information 11-18
Configuring IGMP Filtering 11-20
Default IGMP Filtering Configuration 11-20
Configuring IGMP Profiles 11-20
Applying IGMP Profiles 11-22
Setting the Maximum Number of IGMP Groups 11-23
Displaying IGMP Filtering Configuration 11-24

CHAPTER 12 **Configuring Port-Based Traffic Control 12-1**

Configuring Storm Control 12-1
Understanding Storm Control 12-1
Default Storm Control Configuration 12-3
Enabling Storm Control 12-3
Disabling Storm Control 12-4
Configuring Protected Ports 12-5

Configuring Port Blocking 12-6
Blocking Flooded Traffic on an Interface 12-6
Resuming Normal Forwarding on a Port 12-7
Configuring Port Security 12-8
Understanding Port Security 12-8
Default Port Security Configuration 12-9
Configuration Guidelines 12-9
Enabling and Configuring Port Security 12-9
Displaying Port-Based Traffic Control Settings 12-11

CHAPTER 13 **Configuring CDP 13-1**
Understanding CDP 13-1
Configuring CDP 13-2
Default CDP Configuration 13-2
Configuring the CDP Characteristics 13-2
Disabling and Enabling CDP 13-3
Disabling and Enabling CDP on an Interface 13-4
Monitoring and Maintaining CDP 13-5

CHAPTER 14 **Configuring UDLD 14-1**
Understanding UDLD 14-1
Configuring UDLD 14-3
Default UDLD Configuration 14-3
Enabling UDLD Globally 14-3
Enabling UDLD on an Interface 14-4
Resetting an Interface Shut Down by UDLD 14-4
Displaying UDLD Status 14-5

CHAPTER 15 **Configuring SPAN 15-1**
Understanding SPAN 15-1
SPAN Concepts and Terminology 15-2
SPAN Session 15-2
Traffic Types 15-3
Source Port 15-4
Destination Port 15-4
VLAN-Based SPAN 15-5
SPAN Traffic 15-5
SPAN Interaction with Other Features 15-5

Configuring SPAN 15-6
Default SPAN Configuration 15-7
SPAN Configuration Guidelines 15-7
Creating a SPAN Session and Specifying Ports to Monitor 15-8
Removing Ports from a SPAN Session 15-10
Specifying VLANs to Monitor 15-11
Specifying VLANs to Filter 15-12
Displaying SPAN Status 15-13

CHAPTER 16 **Configuring RMON** 16-1
Understanding RMON 16-1
Configuring RMON 16-2
Default RMON Configuration 16-3
Configuring RMON Alarms and Events 16-3
Configuring RMON Collection on an Interface 16-5
Displaying RMON Status 16-6

CHAPTER 17 **Configuring System Message Logging** 17-1
Understanding System Message Logging 17-1
Configuring System Message Logging 17-2
System Log Message Format 17-2
Default System Message Logging Configuration 17-3
Disabling and Enabling Message Logging 17-4
Setting the Message Display Destination Device 17-4
Synchronizing Log Messages 17-6
Enabling and Disabling Timestamps on Log Messages 17-7
Enabling and Disabling Sequence Numbers in Log Messages 17-8
Defining the Message Severity Level 17-8
Limiting Syslog Messages Sent to the History Table and to SNMP 17-10
Configuring UNIX Syslog Servers 17-10
Logging Messages to a UNIX Syslog Daemon 17-11
Configuring the UNIX System Logging Facility 17-11
Displaying the Logging Configuration 17-12

CHAPTER 18 **Configuring SNMP** 18-1

Understanding SNMP 18-1
SNMP Versions 18-2
SNMP Manager Functions 18-2
SNMP Agent Functions 18-3
SNMP Community Strings 18-3
Using SNMP to Access MIB Variables 18-3
Configuring SNMP 18-4
Default SNMP Configuration 18-4
Disabling the SNMP Agent 18-5
Configuring Community Strings 18-5
Configuring Trap Managers and Enabling Traps 18-7
Setting the Agent Contact and Location Information 18-9
Limiting TFTP Servers Used Through SNMP 18-9
SNMP Examples 18-10
Displaying SNMP Status 18-10

CHAPTER 19 **Configuring Network Security with ACLs** 19-1

Understanding ACLs 19-1
Supported ACLs 19-2
Router ACLs 19-2
VLAN Maps 19-3
Handling Fragmented and Unfragmented Traffic 19-4
Configuring Router ACLs 19-5
Hardware and Software Handling of Router ACLs 19-5
Unsupported Features 19-6
Creating Standard and Extended IP ACLs 19-6
Access List Numbers 19-7
Creating a Numbered Standard ACL 19-8
Creating a Numbered Extended ACL 19-9
Creating Named Standard and Extended ACLs 19-14
Applying Time Ranges to ACLs 19-15
Including Comments About Entries in ACLs 19-18
Applying the ACL to an Interface or Terminal Line 19-18
Displaying ACLs and Access Groups 19-20
ACL Configuration Examples 19-22
Numbered ACLs 19-24
Extended ACLs 19-24
Named ACLs 19-24

Time Range Applied to an IP ACL 19-25
Commented IP ACL Entries 19-25
ACL Logging 19-26
Configuring VLAN Maps 19-27
VLAN Map Configuration Guidelines 19-28
Creating Named MAC Extended ACLs 19-28
Creating a VLAN Map 19-30
Examples of ACLs and VLAN Maps 19-30
Applying a VLAN Map to a VLAN 19-32
Displaying VLAN Map Information 19-33
Using VLAN Maps in Your Network 19-33
Wiring Closet Configuration 19-34
Denying Access to a Server on Another VLAN 19-35
Using VLAN Maps with Router ACLs 19-36
Guidelines 19-36
Determining if the ACL Configuration Fits in Hardware 19-37
Examples of Router ACLs and VLAN Maps Applied to VLANs 19-39
ACLs and Switched Packets 19-39
ACLs and Bridged Packets 19-40
ACLs and Routed Packets 19-41
ACLs and Multicast Packets 19-42

CHAPTER 20 **Configuring QoS 20-1**

Understanding QoS 20-1
Basic QoS Model 20-3
Classification 20-4
Classification Based on QoS ACLs 20-7
Classification Based on Class Maps and Policy Maps 20-7
Policing and Marking 20-8
Mapping Tables 20-11
Queueing and Scheduling 20-12
Queueing and Scheduling on Gigabit-Capable Ports 20-12
Queueing and Scheduling on 10/100 Ethernet Ports 20-15
Packet Modification 20-17
Configuring QoS 20-18
Default QoS Configuration 20-18
Configuration Guidelines 20-20
Enabling QoS Globally 20-21
Configuring Classification Using Port Trust States 20-21

Configuring the Trust State on Ports within the QoS Domain 20-22
Configuring the CoS Value for an Interface 20-24
Configuring the DSCP Trust State on a Port Bordering Another QoS Domain 20-25
Configuring a QoS Policy 20-26
Classifying Traffic by Using ACLs 20-27
Classifying Traffic by Using Class Maps 20-30
Classifying, Policing, and Marking Traffic by Using Policy Maps 20-32
Classifying, Policing, and Marking Traffic by Using Aggregate Policers 20-37
Configuring DSCP Maps 20-39
Configuring the CoS-to-DSCP Map 20-39
Configuring the IP-Precedence-to-DSCP Map 20-40
Configuring the Policed-DSCP Map 20-41
Configuring the DSCP-to-CoS Map 20-42
Configuring the DSCP-to-DSCP-Mutation Map 20-43
Configuring Egress Queues on Gigabit-Capable Ethernet Ports 20-44
Mapping CoS Values to Select Egress Queues 20-45
Configuring the Egress Queue Size Ratios 20-46
Configuring Tail-Drop Threshold Percentages 20-47
Configuring WRED Drop Thresholds Percentages 20-48
Configuring the Egress Expedite Queue 20-50
Allocating Bandwidth among Egress Queues 20-50
Configuring Egress Queues on 10/100 Ethernet Ports 20-51
Mapping CoS Values to Select Egress Queues 20-52
Configuring the Minimum-Reserve Levels 20-53
Configuring the Egress Expedite Queue 20-54
Allocating Bandwidth among Egress Queues 20-54
Displaying QoS Information 20-56
QoS Configuration Examples 20-56
QoS Configuration for the Common Wiring Closet 20-57
QoS Configuration for the Intelligent Wiring Closet 20-58
QoS Configuration for the Distribution Layer 20-59

CHAPTER 21 **Configuring EtherChannel** 21-1
Understanding EtherChannel 21-1
Understanding Port-Channel Interfaces 21-2
Understanding the Port Aggregation Protocol 21-3
PAgP Modes 21-4
Physical Learners and Aggregate-Port Learners 21-5
PAgP Interaction with Other Features 21-5

Understanding Load Balancing and Forwarding Methods 21-5
Configuring EtherChannel 21-7
Default EtherChannel Configuration 21-7
EtherChannel Configuration Guidelines 21-8
Configuring Layer 2 EtherChannels 21-9
Configuring Layer 3 EtherChannels 21-11
Creating Port-Channel Logical Interfaces 21-11
Configuring the Physical Interfaces 21-12
Configuring EtherChannel Load Balancing 21-13
Configuring the PAgP Learn Method and Priority 21-14
Displaying EtherChannel and PAgP Status 21-16

CHAPTER 22 **Configuring IP Unicast Routing** 22-1
Understanding Routing 22-2
Steps for Configuring Routing 22-3
Configuring IP Addressing 22-4
Default Addressing Configuration 22-4
Assigning IP Addresses to Network Interfaces 22-5
Use of Subnet Zero 22-8
Classless Routing 22-8
Configuring Address Resolution Methods 22-10
Define a Static ARP Cache 22-11
Set ARP Encapsulation 22-12
Enable Proxy ARP 22-13
Routing Assistance When IP Routing is Disabled 22-14
Proxy ARP 22-14
Default Gateway 22-15
ICMP Router Discovery Protocol (IRDP) 22-15
Configuring Broadcast Packet Handling 22-17
Enabling Directed Broadcast-to-Physical Broadcast Translation 22-17
Forwarding UDP Broadcast Packets and Protocols 22-18
Establishing an IP Broadcast Address 22-20
Flooding IP Broadcasts 22-20
Monitoring and Maintaining IP Addressing 22-21
Enabling IP Routing 22-24
Configuring RIP 22-25
RIP Authentication 22-28
Summary Addresses and Split Horizon 22-28

Configuring IGRP 22-30
Load Balancing and Traffic Distribution Control 22-31
Split Horizon 22-34
Configuring OSPF 22-35
OSPF Interface Parameters 22-38
OSPF Area Parameters 22-39
Other OSPF Behavior Parameters 22-41
Change LSA Group Pacing 22-43
Loopback Interface 22-43
Monitoring OSPF 22-44
Configuring EIGRP 22-46
EIGRP Router Mode Commands 22-48
EIGRP Interface Mode Commands 22-49
Configure EIGRP Route Authentication 22-50
Monitoring and Maintaining EIGRP 22-51
Configuring Protocol-Independent Features 22-53
Configuring Cisco Express Forwarding 22-53
Configuring the Number of Equal-Cost Routing Paths 22-54
Configuring Static Routes 22-55
Specifying Default Routes 22-56
Specifying a Default Network 22-56
Redistributing Routing Information 22-57
Filtering Routing Information 22-61
Setting Passive Interfaces 22-61
Controlling Advertising and Processing in Routing Updates 22-62
Filtering Sources of Routing Information 22-62
Managing Authentication Keys 22-63
Monitoring and Maintaining the IP Network 22-64

CHAPTER 23 **Configuring HSRP 23-1**

Understanding HSRP 23-1
Configuring HSRP 23-3
Default HSRP Configuration 23-4
Enabling HSRP 23-4
Configuring HSRP Group Attributes 23-6
Configuring HSRP Priority 23-6
Configuring HSRP Authentication and Timers 23-8
Configuring HSRP Groups and Clustering 23-9
Displaying HSRP Configurations 23-10

CHAPTER 24 **Configuring IP Multicast Routing** 24-1

Cisco Implementation of IP Multicast Routing 24-2
Understanding IGMP 24-3
IGMP Version 1 24-3
IGMP Version 2 24-4
Understanding PIM 24-5
PIM Versions 24-5
PIM Modes 24-5
Auto-RP 24-8
Bootstrap Router 24-8
Multicast Forwarding and Reverse Path Check 24-9
Neighbor Discovery 24-10
Understanding DVMRP 24-11
DVMRP Neighbor Discovery 24-11
DVMRP Route Table 24-11
DVMRP Source Distribution Tree 24-11
Understanding CGMP 24-11
Joining a Group with CGMP 24-12
Leaving a Group with CGMP 24-13
Configuring IP Multicast Routing 24-13
Default Multicast Routing Configuration 24-13
Multicast Routing Configuration Guidelines 24-14
PIMv1 and PIMv2 Interoperability 24-14
Auto-RP and BSR Configuration Guidelines 24-15
Configuring Basic Multicast Routing 24-15
Configuring a Rendezvous Point 24-17
Manually Assigning an RP to Multicast Groups 24-17
Configuring Auto-RP 24-18
Configuring PIMv2 BSR 24-22
Using Auto-RP and a BSR 24-27
Monitoring the RP Mapping Information 24-27
Troubleshooting PIMv1 and PIMv2 Interoperability Problems 24-28
Configuring Advanced PIM Features 24-28
Understanding PIM Shared Tree and Source Tree 24-28
Delaying the Use of PIM Shortest-Path Tree 24-29
Modifying the PIM Router-Query Message Interval 24-30
Configuring Optional IGMP Features 24-31
Default IGMP Configuration 24-31
Changing the IGMP Version 24-32

Changing the IGMP Query Timeout for IGMPv2 24-32
Changing the Maximum Query Response Time for IGMPv2 24-33
Configuring the Multilayer Switch as a Member of a Group 24-34
Controlling Access to IP Multicast Groups 24-35
Modifying the IGMP Host-Query Message Interval 24-36
Configuring the Multilayer Switch as a Statically Connected Member 24-36
Configuring Optional Multicast Routing Features 24-37
Enabling CGMP Server Support 24-38
Configuring sdr Listener Support 24-39
Enabling sdr Listener Support 24-39
Limiting How Long an sdr Cache Entry Exists 24-39
Configuring the TTL Threshold 24-40
Configuring an IP Multicast Boundary 24-42
Configuring Basic DVMRP Interoperability Features 24-43
Configuring DVMRP Interoperability 24-44
Controlling Unicast Route Advertisements 24-44
Configuring a DVMRP Tunnel 24-46
Advertising Network 0.0.0.0 to DVMRP Neighbors 24-48
Responding to mrinfo Requests 24-49
Configuring Advanced DVMRP Interoperability Features 24-50
Enabling DVMRP Unicast Routing 24-50
Rejecting a DVMRP Nonpruning Neighbor 24-51
Controlling Route Exchanges 24-53
Limiting the Number of DVMRP Routes Advertised 24-53
Changing the DVMRP Route Threshold 24-54
Configuring a DVMRP Summary Address 24-54
Disabling DVMRP Autosummarization 24-56
Adding a Metric Offset to the DVMRP Route 24-56
Monitoring and Maintaining IP Multicast Routing 24-57
Clearing Caches, Tables, and Databases 24-58
Displaying System and Network Statistics 24-58
Monitoring IP Multicast Routing 24-59

CHAPTER 25 **Configuring MSDP** 25-1
Understanding MSDP 25-1
MSDP Operation 25-2
MSDP Benefits 25-3
Configuring MSDP 25-4
Default MSDP Configuration 25-4

Configuring a Default MSDP Peer 25-4
Caching Source-Active State 25-6
Requesting Source Information from an MSDP Peer 25-8
Controlling Source Information that Your Switch Originates 25-8
Redistributing Sources 25-9
Filtering Source-Active Request Messages 25-11
Controlling Source Information that Your Switch Forwards 25-12
Using a Filter 25-12
Using TTL to Limit the Multicast Data Sent in SA Messages 25-14
Controlling Source Information that Your Switch Receives 25-14
Configuring an MSDP Mesh Group 25-16
Shutting Down an MSDP Peer 25-16
Including a Bordering PIM Dense-Mode Region in MSDP 25-17
Configuring an Originating Address other than the RP Address 25-18
Monitoring and Maintaining MSDP 25-19

CHAPTER 26 **Configuring Fallback Bridging** 26-1
Understanding Fallback Bridging 26-1
Configuring Fallback Bridging 26-3
Default Fallback Bridging Configuration 26-3
Creating a Bridge Group 26-4
Preventing the Forwarding of Dynamically Learned Stations 26-5
Configuring the Bridge Table Aging Time 26-6
Filtering Frames by a Specific MAC Address 26-6
Adjusting Spanning-Tree Parameters 26-7
Changing the Switch Priority 26-8
Changing the Interface Priority 26-8
Assigning a Path Cost 26-9
Adjusting BPDU Intervals 26-10
Disabling the Spanning Tree on an Interface 26-12
Monitoring and Maintaining the Network 26-12

CHAPTER 27 **Troubleshooting** 27-1
Using Recovery Procedures 27-1
Recovering from Corrupted Software 27-2
Recovering from a Lost or Forgotten Password 27-3
Password Recovery with Password Recovery Enabled 27-3
Procedure with Password Recovery Disabled 27-5

Recovering from a Command Switch Failure 27-7
Replacing a Failed Command Switch with a Cluster Member 27-7
Replacing a Failed Command Switch with Another Switch 27-9
Recovering from Lost Member Connectivity 27-10
Preventing Autonegotiation Mismatches 27-10
Diagnosing Connectivity Problems 27-11
Understanding Ping 27-11
Executing Ping 27-11
Understanding IP Traceroute 27-12
Executing IP Traceroute 27-13
Using Debug Commands 27-14
Enabling Debugging on a Specific Feature 27-14
Enabling All-System Diagnostics 27-15
Redirecting Debug and Error Message Output 27-15
Using the show forward Command 27-15
Using the crashinfo File 27-17

APPENDIX A **Supported MIBs** A-1
MIB List A-1
Using FTP to Access the MIB Files A-2

APPENDIX B **Working with the IOS File System, Configuration Files, and Software Images** B-1
Working with the Flash File System B-1
Displaying Available File Systems B-2
Setting the Default File System B-3
Displaying Information about Files on a File System B-3
Changing Directories and Displaying the Working Directory B-3
Creating and Removing Directories B-4
Copying Files B-4
Deleting Files B-5
Creating, Displaying, and Extracting tar Files B-6
Creating a tar File B-6
Displaying the Contents of a tar File B-6
Extracting a tar File B-7
Displaying the Contents of a File B-8
Working with Configuration Files B-8
Guidelines for Creating and Using Configuration Files B-9
Configuration File Types and Location B-9

Creating a Configuration File By Using a Text Editor B-10
Copying Configuration Files By Using TFTP B-10
Preparing to Download or Upload a Configuration File By Using TFTP B-10
Downloading the Configuration File By Using TFTP B-11
Uploading the Configuration File By Using TFTP B-12
Copying Configuration Files By Using FTP B-12
Preparing to Download or Upload a Configuration File By Using FTP B-13
Downloading a Configuration File By Using FTP B-13
Uploading a Configuration File By Using FTP B-15
Copying Configuration Files By Using RCP B-16
Preparing to Download or Upload a Configuration File By Using RCP B-16
Downloading a Configuration File By Using RCP B-17
Uploading a Configuration File By Using RCP B-18
Clearing Configuration Information B-19
Clearing the Startup Configuration File B-19
Deleting a Stored Configuration File B-19
Working with Software Images B-19
Image Location on the Switch B-20
tar File Format of Images on a Server or Cisco.com B-20
Copying Image Files By Using TFTP B-21
Preparing to Download or Upload an Image File By Using TFTP B-22
Downloading an Image File By Using TFTP B-22
Uploading an Image File By Using TFTP B-24
Copying Image Files By Using FTP B-25
Preparing to Download or Upload an Image File By Using FTP B-25
Downloading an Image File By Using FTP B-26
Uploading an Image File By Using FTP B-28
Copying Image Files By Using RCP B-29
Preparing to Download or Upload an Image File By Using RCP B-29
Downloading an Image File By Using RCP B-30
Uploading an Image File By Using RCP B-32

APPENDIX C **Unsupported CLI Commands** C-1
Access Control Lists C-1
Unsupported Privileged EXEC Commands C-1
ARP Commands C-1
Unsupported Global Configuration Commands C-1
Unsupported Interface Configuration Commands C-1

FallBack Bridging C-2
Unsupported Privileged EXEC Commands C-2
Unsupported Global Configuration Commands C-2
Unsupported Interface Configuration Commands C-2
HSRP C-3
Unsupported Global Configuration Commands C-3
Unsupported Interface Configuration Commands C-3
Interface Configuration Commands C-4
IP Multicast Routing C-4
Unsupported Privileged EXEC Commands C-4
Unsupported Global Configuration Commands C-4
Unsupported Interface Configuration Commands C-5
IP Unicast Routing C-5
Unsupported Privileged EXEC or User EXEC Commands C-5
Unsupported Global Configuration Commands C-5
Unsupported Interface Configuration Commands C-6
Unsupported VPN Configuration Commands C-6
Unsupported VRF Configuration Commands C-6
Unsupported Route Map Commands C-6
MSDP C-7
Unsupported Privileged EXEC Commands C-7
Unsupported Global Configuration Commands C-7
RADIUS C-7
Unsupported Global Configuration Commands C-7

INDEX

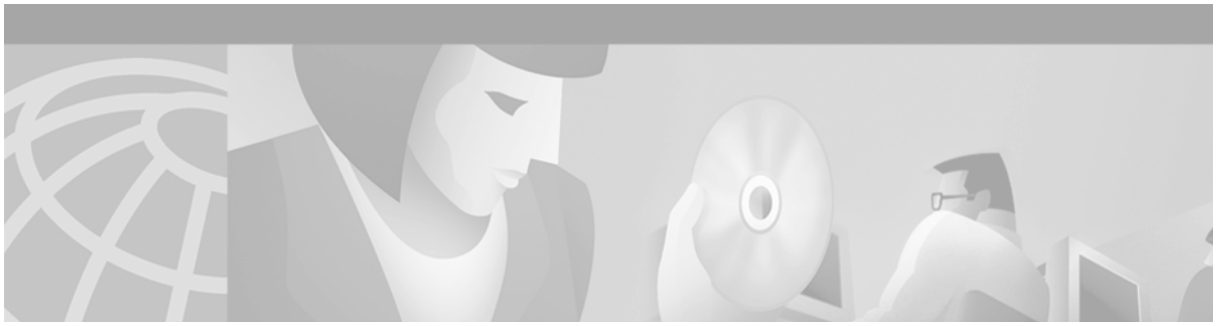

Preface

Audience

This guide is for the networking professional managing the Catalyst 3550 switch, hereafter referred to as the switch or the multilayer switch. Before using this guide, you should have experience working with the Cisco IOS and be familiar with the concepts and terminology of Ethernet and local area networking.

Purpose

This guide provides the information you need to configure Layer 2 and Layer 3 software features on your switch. The Catalyst 3550 switch is supported by either the standard multilayer software image (SMI) or the enhanced multilayer software image (EMI). The EMI provides a richer set of enterprise-class features, including hardware-based IP unicast and multicast routing, inter-VLAN routing, routed access control lists (ACLs), and the Hot Standby Router Protocol (HSRP). All Catalyst 3550 Gigabit Ethernet switches are shipped with the EMI pre-installed. Catalyst 3550 Fast Ethernet switches are shipped with either the SMI or the EMI pre-installed. After initial deployment, you can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from running the SMI to the EMI.

This guide provides procedures for using the commands that have been created or changed for use with the Catalyst 3550 switch. It does not provide detailed information about these commands. For detailed information about these commands, refer to the C*atalyst 3550 Multilayer Switch Command Reference* for this release. For information about the standard IOS Release 12.1 commands, refer to the IOS documentation set available from the Cisco.com home page at **Service and Support > Technical Documents**. On the Cisco Product Documentation home page, select **Release 12.1** from the Cisco IOS Software drop-down list.

This guide also includes an overview of the Cluster Management Suite (CMS) web-based, switch management interface, which helps you create and manage clusters of switches. This guide does not provide field-level descriptions of the CMS windows nor does it provide the procedures for configuring switches and switch clusters from CMS. For all CMS window descriptions and procedures, refer to the CMS online help, which is integrated with the software image.

This guide does not describe system messages you might encounter or how to install your switch. For more information, refer to the *Catalyst 3550 Multilayer Switch System Message Guide* for this release and to the *Catalyst 3550 Multilayer Switch Hardware Installation Guide*.

Organization

This guide is organized into these chapters:

Chapter 1, "Overview," lists the software features of this release and provides examples of how the switch can be deployed in a network.

Chapter 2, "Using the Command-Line Interface," describes how to access the command modes, use the command-line interface (CLI), and describes CLI messages that you might receive. It also describes how to get help, abbreviate commands, use **no** and **default** forms of commands, use command history and editing features, and how to search and filter the output of **show** and **more** commands.

Chapter 3, "Getting Started with CMS," describes the Cluster Management Suite (CMS) web-based, switch management interface. For information on configuring your web browser and accessing CMS, refer to the release notes. For field-level descriptions of all CMS windows and procedures for using the CMS windows, refer to the online help.

Chapter 4, "Assigning the Switch IP Address and Default Gateway," describes how to create the initial switch configuration (for example, assign the switch IP address and default gateway information) by using a variety of automatic and manual methods. It also describes how to modify the switch startup configuration.

Chapter 5, "Clustering Switches," describes switch clusters and the considerations for creating and maintaining them. The online help provides the CMS procedures for configuring switch clusters. Configuring switch clusters is most easily performed through CMS; therefore, CLI procedures are not provided. Cluster commands are described in the *Catalyst 3550 Multilayer Switch Command Reference*.

Chapter 6, "Administering the Switch," describes how to perform one-time operations to administer your switch. It describes how to prevent unauthorized access to your switch through the use of passwords, privilege levels, the Terminal Access Controller Access Control System Plus (TACACS+), and the Remote Authentication Dial-In User Service (RADIUS). It also describes how to set the system date and time, system name and prompt, create a login banner, how to manage the MAC address table, and how to optimize system resources for user-selected features.

Chapter 7, "Configuring 802.1X Port-Based Authentication," describes how to configure 802.1X port-based authentication to prevent unauthorized devices (clients) from gaining access to the network. As LANs extend to hotels, airports, and corporate lobbies, insecure environments could be created.

Chapter 8, "Configuring Interface Characteristics," defines the types of Layer 2 and Layer 3 interfaces on the switch. It describes the **interface** command and provides procedures for configuring physical interfaces.

Chapter 9, "Creating and Maintaining VLANs," describes how to create and maintain VLANs. It includes information about VLAN modes, the VLAN Trunking Protocol (VTP) database, and the VLAN Membership Policy Server (VMPS).

Chapter 10, "Configuring STP," describes how to configure basic and advanced spanning-tree features.

Chapter 11, "Configuring IGMP Snooping and MVR," describes how to configure Layer 2 Internet Group Management Protocol (IGMP) snooping. It also describes Multicast VLAN Registration (MVR), a local IGMP snooping feature available on the switch, and how to use IGMP filtering to control multicast group membership.

Chapter 12, "Configuring Port-Based Traffic Control," describes how to reduce traffic storms by setting broadcast, multicast, and unicast storm-control threshold levels; how to protect ports from receiving traffic from other ports on a switch; how to block unknown broadcast and unicast traffic; and how to configure port security using secure MAC addresses.

Chapter 13, "Configuring CDP," describes how to configure Cisco Discovery Protocol (CDP) on your switch.

Chapter 14, “Configuring UDLD,” describes how to configure the UniDirectional Link Detection (UDLD) feature. UDLD enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists.

Chapter 15, “Configuring SPAN,” describes how to configure Switch Port Analyzer (SPAN), which selects network traffic for analysis by a network analyzer such as a SwitchProbe device or other Remote Monitoring (RMON) probe. SPAN mirrors traffic received or sent (or both) on a source port, or traffic received on one or more source ports or source VLANs, to a destination port.

Chapter 16, “Configuring RMON,” describes how to configure remote monitoring (RMON). The RMON feature, which is used with the Simple Network Management Protocol (SNMP) agent in the switch, means that you can monitor all the traffic flowing among switches on all connected LAN segments.

Chapter 17, “Configuring System Message Logging,” describes how to configure system message logging. It describes the message format, how to change the message display destination device, limit the type of messages sent, configure UNIX server syslog daemon and define the UNIX system logging facility, and timestamp messages.

Chapter 18, “Configuring SNMP,” describes how to configure the Simple Network Management Protocol (SNMP). It describes how to configure community strings, enable trap managers and traps, set the agent contact and location information, and how to limit TFTP servers used through SNMP.

Chapter 19, “Configuring Network Security with ACLs,” describes how to configure network security on your switch using two types of access control lists (ACLs), router ACLs and VLAN maps. It describes how to apply ACLs to interfaces and provides examples.

Chapter 20, “Configuring QoS,” describes how to configure quality of service (QoS) on your switch. With this feature, you can provide preferential treatment to certain types traffic.

Chapter 21, “Configuring EtherChannel,” describes how to bundle a set of individual ports into a single logical link on Layer 2 and Layer 3 interfaces.

Chapter 22, “Configuring IP Unicast Routing,” describes how to configuring IP unicast routing on your switch, including configuring IP addressing features, Routing Information Protocol (RIP), Interior Gateway Routing Protocol (IGRP), Open Shortest Path First (OSPF) protocol, and Enhanced IGRP (EIGRP). To use this feature, you must have the enhanced multilayer software image installed on your switch.

Chapter 23, “Configuring HSRP,” describes how to use Hot Standby Router Protocol (HSRP) to provide routing redundancy for routing IP traffic without depending on the availability of any single router. To use this feature, you must have the enhanced multilayer software image installed on your switch.

Chapter 24, “Configuring IP Multicast Routing,” how to configuring IP multicast routing. It describes how to use and configure the Internet Group Management Protocol (IGMP), Protocol-Independent Multicast (PIM) protocol, Cisco Group Management Protocol (CGMP) server functionality, and how to inter-operate between PIM and Distance Vector Multicast Routing Protocol (DVMRP) domains. To use this feature, you must have the enhanced multilayer software image installed on your switch.

Chapter 25, “Configuring MSDP,” describes how to configure the Multicast Source Discovery Protocol (MSDP), which is a mechanism to connect multiple PIM sparse-mode domains. To use this feature, you must have the enhanced multilayer software image installed on your switch.

Chapter 26, “Configuring Fallback Bridging,” describes how to configure fallback bridging on your switch. With fallback bridging, you can bridge non-IP protocols between VLAN bridge domains and routed ports. To use this feature, you must have the enhanced multilayer software image installed on your switch

Chapter 27, “Troubleshooting,” describes how to identify and resolve software problems related to the IOS software.

Appendix A, "Supported MIBs," lists the supported MIBs for this release and how to use FTP to access the MIB files.

Appendix B, "Working with the IOS File System, Configuration Files, and Software Images," describes how to manipulate the Flash file system, how to copy configuration files, and how to archive (upload and download) software images.

Appendix C, "Unsupported CLI Commands," lists the unsupported command-line interface (CLI) commands that are displayed when you enter the question mark (?) at the switch prompt. The unsupported commands are listed by software feature and command mode.

Conventions

This publication uses these conventions to convey instructions and information:

Command descriptions use these conventions:

- Commands and keywords are in **boldface** text.
- Arguments for which you supply values are in *italic*.
- Square brackets ([ ]) mean optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate the alternative elements.
- Braces and vertical bars within square brackets ([{ | }]) mean a required choice within an optional element.

Interactive examples use these conventions:

- Terminal sessions and system displays are in `screen` font.
- Information you enter is in **`boldface screen`** font.
- Nonprinting characters, such as passwords or tabs, are in angle brackets (< >).

Notes, cautions, and timesavers use these conventions and symbols:

Note Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this manual.

Caution Means *reader be careful*. In this situation, you might do something that could result equipment damage or loss of data.

Timesaver Means the following *will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information.

Related Publications

These documents provide complete information about the switch and are available from this Cisco.com site:

http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/index.htm

You can order printed copies of documents with a DOC-xxxxxx= number from the Cisco.com sites and from the telephone numbers listed in the "Ordering Documentation" section on page xxxii.

- *Release Notes for the Catalyst 3550 Multilayer Switch* (not orderable but available on Cisco.com)

Note Switch requirements and procedures for initial configurations and software upgrades tend to change and therefore appear only in the release notes. Before installing, configuring, or upgrading the switch, refer to the release notes on Cisco.com for the latest information.

- *Catalyst 3550 Multilayer Switch Software Configuration Guide* (order number DOC-7811194=)
- *Catalyst 3550 Multilayer Switch Command Reference* (order number DOC-7811195=)
- *Catalyst 3550 Multilayer Switch System Message Guide* (order number DOC-7811196=)
- Cluster Management Suite (CMS) online help (available only from the switch CMS software)
- *Catalyst 3550 Multilayer Switch Hardware Installation Guide* (order number DOC-7811358=)
- *1000BASE-T Gigabit Interface Converter Installation Note* (not orderable but is available on Cisco.com)
- *Catalyst GigaStack Gigabit Interface Converter Hardware Installation Guide* (order number DOC-786460=)

Obtaining Documentation

The following sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following URL:

http://www.cisco.com

Translated documentation is available at the following URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco Direct Customers can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can send us your comments by completing the online survey. When you display the document listing for this platform, click **Give Us Your Feedback**. After you display the survey, select the manual that you wish to comment on. Click **Submit** to send your comments to the Cisco documentation group.

You can e-mail your comments to bug-doc@cisco.com.

To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Cisco Systems, Inc.
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you to

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

You can self-register on Cisco.com to obtain customized information and service. To access Cisco.com, go to the following URL:

http://www.cisco.com

Technical Assistance Center

The Cisco TAC is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two types of support are available through the Cisco TAC: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Inquiries to Cisco TAC are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

Which Cisco TAC resource you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

The Cisco TAC Web Site allows you to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to the following URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco services contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to the following URL to register:

http://www.cisco.com/register/

If you cannot resolve your technical issues by using the Cisco TAC Web Site, and you are a Cisco.com registered user, you can open a case online by using the TAC Case Open tool at the following URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, it is recommended that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses issues that are classified as priority level 1 or priority level 2; these classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer will automatically open a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to the following URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled; for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). In addition, please have available your service agreement number and your product serial number.

Overview

This chapter provides these topics about the Catalyst 3550 multilayer switch software:

- Features, page 1-1
- Management Options, page 1-5
- Network Configuration Examples, page 1-7

Features

The Catalyst 3550 software supports the hardware listed in the release notes. These sections describe the features supported in this release.

Note Layer 3 (routing) features require that you have the enhanced multilayer software image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from running the SMI to the EMI.

Table 1-1 Features

Ease of Use and Ease of Deployment
• Cluster Management Suite (CMS) software for simplifying switch and switch cluster management through a web browser, such as Netscape Communicator or Microsoft Internet Explorer, from anywhere in your intranet • Switch clustering technology used with CMS, for – Unified configuration, monitoring, authentication, and software upgrade of multiple switches (refer to the release notes for a list of eligible cluster members). – Automatic discovery of candidate switches and creation of clusters of up to 16 switches that can be managed through a single IP address. – Extended discovery of cluster candidates that are not directly connected to the command switch. • Hot Standby Router Protocol (HSRP) for command-switch redundancy (requires the enhanced multilayer software image). **Note** See the "Advantages of Using CMS and Clustering Switches" section on page 1-6. Refer to the release notes for the CMS, cluster hardware, software, and browser requirements.

Table 1-1 Features (continued)

Performance
• Autosensing of port speed and autonegotiation of duplex mode on all switch ports for optimizing bandwidth • IEEE 802.3X flow control on all Ethernet ports • EtherChannel for enhanced fault tolerance and for providing up to 8 Gbps (Gigabit EtherChannel) or 800 Mbps (Fast EtherChannel) full duplex of bandwidth between switches, routers, and servers • Port Aggregation Protocol (PAgP) for automatic creation of EtherChannel links • Per-port storm control for preventing broadcast, multicast, and unicast storms • Port blocking on forwarding unknown unicast and multicast traffic • Cisco Group Management Protocol (CGMP) server support and Internet Group Management Protocol (IGMP) snooping for IGMP versions 1 and 2: – (For CGMP devices) CGMP for limiting multicast traffic to specified end stations and reducing overall network traffic – (For IGMP devices) IGMP snooping for limiting flooding of multicast traffic • Multicast VLAN registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security reasons • IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong • System Database Management (SDM) templates for allocating system resources to maximize support for user-selected features
Manageability
• Dynamic Host Configuration Protocol (DHCP) for automating configuration of switch information (such as IP address, default gateway, host name, and Domain Name System [DNS] and Trivial File Transfer Protocol [TFTP] server names) • Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding host name and to a TFTP server for administering software upgrades from a TFTP server • Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding Media Access Control (MAC) address • Cisco Discovery Protocol (CDP) versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network • Network Time Protocol (NTP) for providing a consistent timestamp to all switches from an external source • Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses • In-band management access through CMS • In-band management access through up to 16 simultaneous Telnet connections for multiple command-line interface (CLI)-based sessions over the network • In-band management access through Simple Network Management Protocol (SNMP) versions 1 and 2c get and set requests • Out-of-band management access through the switch console port to a directly attached terminal or to a remote terminal through a serial connection or a modem **Note** For additional descriptions of the management interfaces, see the "Management Options" section on page 1-5.

Table 1-1 Features (continued)

Redundancy
• Hot Standby Router Protocol (HSRP) for command switch and Layer 3 router redundancy • UniDirectional Link Detection (UDLD) on all Ethernet ports for detecting and disabling unidirectional links on fiber-optic interfaces caused by incorrect fiber-optic wiring or port faults • IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features: – Per-VLAN Spanning Tree (PVST) for balancing load across VLANs – Port Fast mode for eliminating forward delay by enabling a port to immediately change from a blocking state to a forwarding state – UplinkFast, cross-stack UplinkFast, and BackboneFast for fast convergence after a spanning-tree topology change and for achieving load balancing between redundant uplinks, including Gigabit uplinks and cross-stack Gigabit uplinks – STP root guard for preventing switches outside the network core from becoming the STP root **Note** The switch supports up to 128 spanning-tree instances.
VLAN Support
• Support for up to 1005 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth • VLAN Query Protocol (VQP) for dynamic VLAN membership • Inter-Switch Link (ISL) and IEEE 802.1Q trunking encapsulation on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources • Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q or ISL) to be used • VLAN Trunking Protocol (VTP) and VTP pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic
Security
• Password-protected access (read-only and read-write access) to management interfaces (CMS and CLI) for protection against unauthorized configuration changes • Multilevel security for a choice of security level, notification, and resulting actions • Static MAC addressing for ensuring security • Protected port option for restricting the forwarding of traffic to designated ports on the same switch • Port security option for limiting and identifying MAC addresses of the stations allowed to access the port • Bridge Protocol Data Unit (BPDU) Guard for shutting down a Port Fast-configured port when an invalid configuration occurs • Standard and extended IP access control lists (ACLs) for defining security policies on routed interfaces • VLAN ACLs (VLAN maps) for providing intra-VLAN security by filtering traffic based on information in the MAC, IP, and TCP/User Datagram Protocol (UDP) headers • Source and destination MAC-based ACLs for filtering non-IP traffic • IEEE 802.1X port-based authentication to prevent unauthorized devices (clients) from gaining access to the network

Table 1-1 Features (continued)

• Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server • Remote Authentication Dial-In User Service (RADIUS), which provides detailed accounting information and flexible administrative control over authentication and authorization processes
Quality of Service and Class of Service
Classification
• IP type-of-service/Differentiated Services Code Point (IP TOS/DSCP) and 802.1P class of service (CoS) marking priorities on a per-port basis for protecting the performance of mission-critical applications • IP TOS/DSCP and 802.1P CoS marking based on flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance quality of service at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network
Policing
• Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow • Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates • Up to 128 policers on ingress Gigabit-capable Ethernet ports Up to eight policers on ingress 10/100 ports Up to eight policers per egress port (aggregate policers only)
Out-of-Profile
• Out-of-profile markdown for packets that exceed bandwidth utilization limits
Egress Policing and Scheduling of Egress Queues
• Four egress queues on all switch ports. These queues can either be configured with the Weighted Round Robin (WRR) scheduling algorithm or configured with one queue as a strict priority queue and the other three queues for WRR. The strict priority queue must be empty before the other three queues are serviced. You can use the strict priority queue for mission-critical and time-sensitive traffic. • Tail drop and Weight Random Early Detection (WRED) techniques for avoiding congestion on Gigabit Ethernet ports; tail drop for congestion avoidance on Fast Ethernet ports
Layer 3 Support
Note These features requires the enhanced multilayer software image. • Hot Standby Router Protocol (HSRP) for Layer 3 router redundancy • IP routing protocols for load balancing and for constructing scalable, routed backbones: – Routing Information Protocol (RIP) versions 1 and 2 – Open Shortest Path First (OSPF) – Interior Gateway Routing Protocol (IGRP) and Enhanced IGRP (EIGRP) • IP routing between VLANs (inter-VLAN routing) for full Layer 3 routing between two or more VLANs, allowing each VLAN to maintain its own autonomous data-link domain • Fallback bridging for forwarding non-IP traffic between two or more VLANs • Static IP routing for manually building a routing table of network path information • Equal-cost routing for load balancing and redundancy

Table 1-1 Features (continued)

• Internet Control Message Protocol (ICMP) and ICMP Router Discovery Protocol (IRDP) for using router advertisement and router solicitation messages to discover the addresses of routers on directly attached subnets • Protocol-Independent Multicast (PIM) for multicast routing within the network, allowing for devices in the network to receive the multicast feed requested and for switches not participating in the multicast to be pruned. Includes support for PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode. • Distance Vector Multicast Routing Protocol (DVMRP) tunnelling for interconnecting two multicast-enabled networks across non-multicast networks • DHCP relay for forwarding UDP broadcasts, including IP address requests, from DHCP clients
Monitoring
• Switch LEDs that provide port- and switch-level status • Switch Port Analyzer (SPAN) for traffic monitoring on any port or VLAN • Four groups (history, statistics, alarms, and events) of embedded remote monitoring (RMON) agents for network monitoring and traffic analysis • Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events • MAC address notification for tracking users on a network by storing the MAC addresses that the switch has learned or removed

Management Options

The Catalyst 3550 switch is designed for plug-and-play operation: you need to configure only basic IP information for the switch and connect it to the other devices in your network. If you have specific network needs, you can configure and monitor the switch—on an individual basis or as part of a switch cluster—through its various management interfaces.

Management Interface Options

You can configure and monitor individual switches and switch clusters by using these interfaces:

- CMS—CMS is a graphical user interface that can be launched from anywhere in your network through a web browser such as Netscape Communicator or Microsoft Internet Explorer. CMS is already installed on the switch. Using CMS, you can configure and monitor a standalone switch, a specific cluster member, or an entire switch cluster. You can also display network topologies to gather link information and to display switch images to modify switch and port settings.

  For more information about CMS, see Chapter 3, “Getting Started with CMS.”

- CLI—The switch IOS CLI software is enhanced to support desktop- and multilayer-switching features. You can configure and monitor the switch and switch cluster members from the CLI. You can access the CLI either by connecting your management station directly to the switch console port or by using Telnet from a remote management station.

  For more information about the CLI, see Chapter 2, “Using the Command-Line Interface.”

- SNMP—SNMP provides a means to monitor and control the switch and switch cluster members. You can manage switch configuration settings, performance, security, and collect statistics by using SNMP management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView.

  You can manage the switch from an SNMP-compatible management station that is running platforms such as HP OpenView or SunNet Manager. The switch supports a comprehensive set of MIB extensions and four RMON groups.

  For more information about using SNMP, see Chapter 18, "Configuring SNMP."

Advantages of Using CMS and Clustering Switches

Using CMS and switch clusters can simplify and minimize your configuration and monitoring tasks. You can use Cisco switch clustering technology to manage up to 16 interconnected, supported Catalyst switches through one IP address. This can conserve IP addresses if you have a limited number of them. CMS is the easiest interface to use and makes switch and switch cluster management accessible to authorized users from any PC on your network.

By using switch clusters and CMS, you can

- Manage and monitor interconnected Catalyst switches (refer to the release notes for a list of supported switches), regardless of their geographic proximity and interconnection media, including Ethernet, Fast Ethernet, Fast EtherChannel, Cisco GigaStack Gigabit Interface Converter (GBIC), Gigabit Ethernet, and Gigabit EtherChannel connections.
- Accomplish multiple configuration tasks from a single CMS window without needing to remember CLI commands to accomplish specific tasks.
- Apply actions from CMS to multiple ports and multiple switches at the same time. Here are some examples of configuring and managing multiple ports and switches:
  - Port configuration such as speed and duplex settings
  - Port and console port security settings
  - NTP, STP, VLAN, and QoS configurations
  - Inventory and statistic reporting and link- and switch-level monitoring and troubleshooting
  - Group software upgrades
- View a topology of interconnected devices to identify existing switch clusters and eligible switches that can join a cluster. You can also use the topology to quickly identify link information between switches.
- Monitor real-time status of a switch or multiple switches from the LEDs on the front-panel images. The system, redundant power system (RPS), and port LED colors on the images are similar to those used on the physical LEDs.
- Use an interactive mode that takes you step-by-step through configuring complex features such as VLANs, ACLs, and QoS.
- Use a wizard that prompts you to provide only minimal required information to configure complex features such as QoS priorities for video traffic, priority levels for data applications, and security.

For more information about CMS, see Chapter 3, "Getting Started with CMS." For more information about switch clusters, see Chapter 5, "Clustering Switches."

Network Configuration Examples

This section provides network configuration concepts and includes examples of using the switch in different network topologies.

Design Concepts

As your network users compete for network bandwidth, it takes longer to send and receive data. When you configure your network, consider the bandwidth required by your network users and the relative priority of the network applications they use.

Table 1-2 describes what can cause network performance to degrade and how you can configure your network to increase the bandwidth available to your network users.

Table 1-2 Increasing Network Performance

Network Demands	Suggested Design Methods
Too many users on a single network segment and a growing number of users accessing the Internet	• Create smaller network segments so that fewer users share the bandwidth, and use VLANs and IP subnets to place the network resources in the same logical network as the users who access those resources most. • Use full-duplex operation between the switch and its connected workstations.
• Increased power of new PCs, workstations, and servers • High bandwidth demand from networked applications (such as e-mail with large attached files) and from bandwidth-intensive applications (such as multimedia)	• Connect global resources—such as servers and routers to which the network users require equal access—directly to the high-speed switch ports so that they have their own high-speed segment. • Use the EtherChannel feature between the switch and its connected servers and routers.

Bandwidth alone is not the only consideration when designing your network. As your network traffic profiles evolve, consider providing network services that can support applications for voice and data integration, multimedia integration, application prioritization, and security. Table 1-3 describes some network demands and how you can meet those demands.

Table 1-3 Providing Network Services

Network Demands	Suggested Design Methods
Efficient bandwidth usage for multimedia applications and guaranteed bandwidth for critical applications	• Use IGMP snooping to efficiently forward multimedia and multicast traffic. • Use other QoS mechanisms such as packet classification, marking, scheduling, and congestion avoidance to classify traffic with the appropriate priority level, thereby providing maximum flexibility and support for mission-critical, unicast, and multicast and multimedia applications. • Use optional IP multicast routing to design networks better suited for multicast traffic. • Use MVR to continuously send multicast streams in a multicast VLAN, but to isolate the streams from subscriber VLANs for bandwidth and security reasons.
High demand on network redundancy to provide *always on* mission-critical applications	• Use HSRP for router redundancy. • Use VLAN trunks, cross-stack UplinkFast, and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
An evolving demand for IP telephony	• Use QoS to prioritize applications such as IP telephony during congestion and to help control both delay and jitter within the network. • Use switches that support at least two queues per port to prioritize voice and data traffic as either high- or low-priority, based on 802.1P/Q. • Use voice VLAN IDs (VVIDs) on the Catalyst 2900 XL and 3500 XL switches to provide separate VLANs for voice traffic.
A growing demand for using existing infrastructure to transport data and voice from a home or office to the Internet or an intranet at higher speeds	Use the Catalyst 2900 LRE XL switches to provide up to 15 Mb of IP connectivity over existing infrastructure, such as existing telephone lines. **Note** Long-Reach Ethernet (LRE) is the technology used in the Catalyst 2900 LRE XL switches. Refer to the Catalyst 2900 XL and 3500 XL documentation set about these switches and the LRE technology.

Figure 1-1 shows three configuration examples of using Catalyst switches to create the following:

- Cost-effective wiring closet—A cost-effective way to connect many users to the wiring closet is to connect a Catalyst switch cluster of up to nine Catalyst 3550 XL switches (or with a mix of Catalyst 3550, Catalyst 2950, Catalyst 3500 XL, and Catalyst 2900 XL switches) through GigaStack GBIC connections. To preserve switch connectivity if one switch in the stack fails, connect the bottom switch to the top switch to create a GigaStack loopback, and enable cross-stack UplinkFast on the cross-stack Gigabit uplinks.

  You can have redundant uplink connections, using Gigabit GBIC modules, from the GigaStack cluster to a Gigabit backbone switch such as the Catalyst 3550-12T or Catalyst 3550-12G switch. You can also create backup paths by using Fast Ethernet, Gigabit, or EtherChannel links. If one of the redundant connections fails, the other can serve as a backup path. You can configure the Catalyst 3550-12T or Catalyst 3550-12G switch as a switch cluster manager to manage stack members through a single IP address. The Catalyst 3550-12T or Catalyst 3550-12G switch can be connected to a Gigabit server through a 1000BASE-T connection.

- High-performance workgroup—For high-speed access to network resources, you can use Catalyst 3550 switches in the access layer to provide Gigabit Ethernet to the desktop. To prevent congestion, use QoS DSCP marking priorities on these switches. For high-speed IP forwarding at the distribution layer, connect the Catalyst 3550 switches in the access layer to a Gigabit multilayer switch (such as the Catalyst 3550 multilayer switch) in the backbone.

  Each switch in this configuration provides users with a dedicated 1-Gbps connection to network resources in the backbone. Compare this with the switches in a GigaStack configuration, where the 1-Gbps connection is shared among the switches in the stack. Using these Gigabit GBIC modules also provides flexibility in media and distance options:

  – 1000BASE-SX GBIC: fiber-optic connections of up to 1804 ft (550 m)
  – 1000BASE-LX/LH GBIC: fiber-optic connections of up to 32,808 ft (10 km)
  – 1000BASE-ZX GBIC: fiber-optic connections of up to 328,084 ft (100 km)
  – 1000BASE-T GBIC: copper connections of up to 328 ft (100 m)

- Redundant Gigabit backbone—Using HSRP, you can create backup paths between two Catalyst 3550 multilayer switches to enhance network reliability and load balancing for different VLANs and subnets. Using HSRP also provides faster network convergence if any network failure occurs. You can connect the Catalyst switches, again in a star configuration, to two Catalyst 3550 multilayer backbone switches. If one of the backbone switches fails, the second backbone switch preserves connectivity between the switches and network resources.

Figure 1-1 Example Configurations

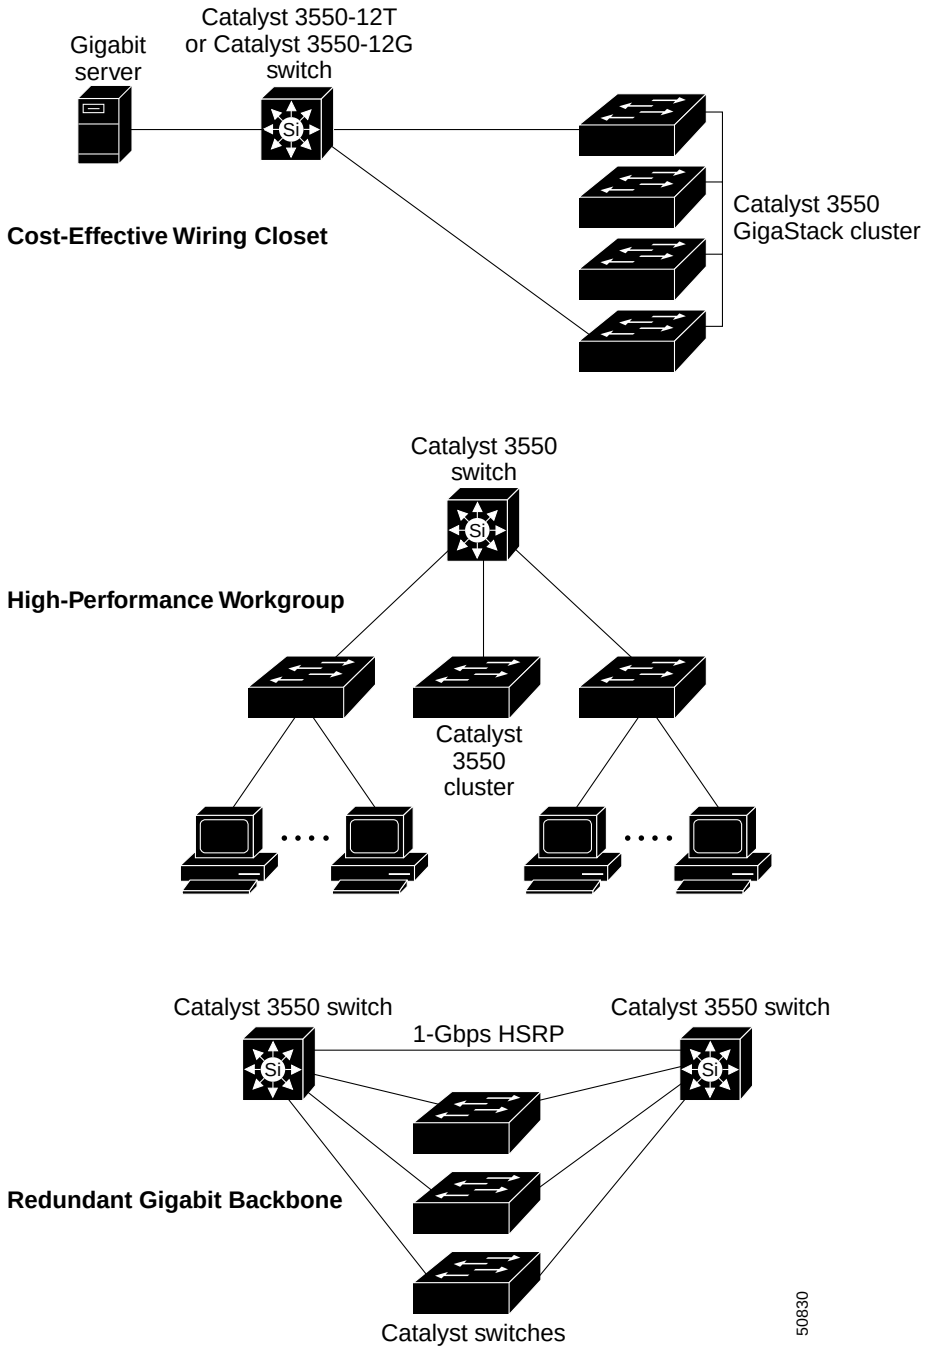

Small to Medium-Sized Network Using Mixed Switches

Figure 1-2 shows a configuration for a network of up to 500 employees. This network uses Catalyst 3550 multilayer switches to aggregate up to ten wiring closets through high-speed uplinks. For network reliability and load balancing, this network includes two routers and two Catalyst 3550 multilayer switches, all with HSRP enabled. This ensures connectivity to the Internet, WAN, and mission-critical network resources if one of the routers or Catalyst 3550 multilayer switches fails.

The wiring closets have a mix of switches such as the Catalyst 3550, Catalyst 3500 XL, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switches. These switches are connected to workstations, Cisco IP Phones, and local servers. You can cluster these switches into multiple clusters, as shown, or into a single cluster. You can manage a cluster through the IP address of its primary and secondary command switches, regardless of the geographic location of the cluster members.

This network uses VLANs to segment the network logically into well-defined broadcast groups and for security management. Data and multimedia traffic are configured on the same VLAN.

When an end station in one VLAN needs to communicate with an end station in another VLAN, a router or multilayer switch routes the traffic to the appropriate destination VLAN. In this network, the Catalyst 3550 multilayer switches provide inter-VLAN routing. VLAN access control lists (VLAN maps) on the Catalyst 3550 switches provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network.

In addition to inter-VLAN routing, the Catalyst 3550 multilayer switches provide QoS mechanisms such as DSCP priorities to prioritize the different types of network traffic and to deliver high-priority traffic in a predictable manner. If congestion occurs, QoS drops low-priority traffic to allow delivery of high-priority traffic.

With the Catalyst 3550 multilayer switches providing inter-VLAN routing and other network services, the routers focus on firewall services, Network Address Translation (NAT) services, voice-over-IP (VoIP) gateway services, and WAN and Internet access.

Figure 1-2 Catalyst 3550 Switches in a Collapsed Backbone Configuration

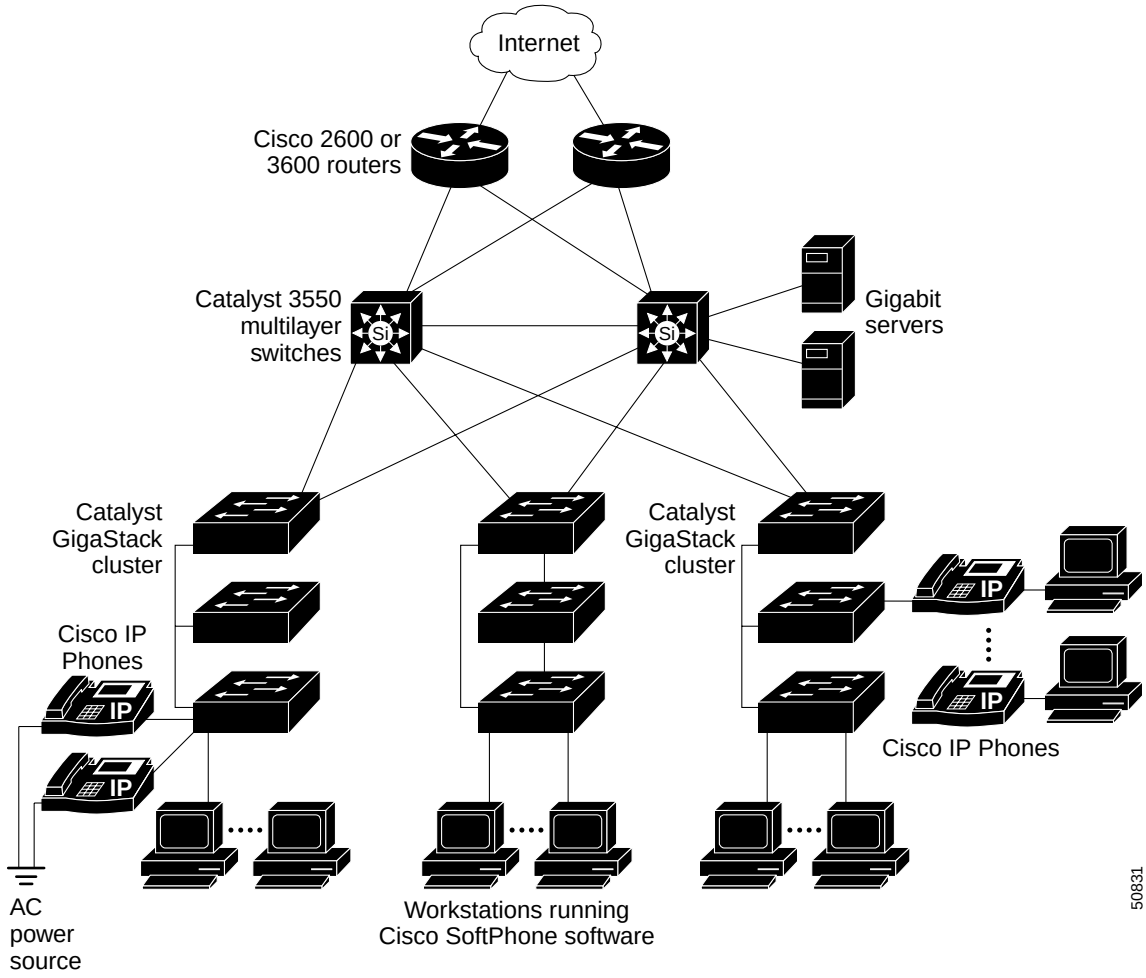

Large Network Using Only Catalyst 3550 Switches

Switches in the wiring closet have traditionally been Layer 2-only devices, but as network traffic profiles evolve, switches in the wiring closet are increasingly employing multilayer services such as multicast management and traffic classification. Figure 1-3 shows a configuration for a network exclusively using Catalyst 3550 multilayer switches in the wiring closets and a Catalyst 6000 switch in the backbone to aggregate up to ten wiring closets.

In the wiring closet, each Catalyst 3550 switch has IGMP snooping enabled to efficiently forward multimedia and multicast traffic. QoS ACLs that either drop or mark nonconforming traffic based on bandwidth limits are also configured on each switch. VLAN maps provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network. QoS features can limit bandwidth on a per-port or per-user basis. The switch ports are configured as either trusted or untrusted. You can configure a trusted port to trust the CoS value, the DSCP value, or the IP precedence. If you configure the port as untrusted, you can use an ACL to mark the frame in accordance with the network policy.

Within each wiring closet is a Catalyst 3550 multilayer switch for inter-VLAN routing. These switches provide proxy ARP services to determine IP and MAC address mapping, thereby removing this task from the routers and lessening this type of traffic on the WAN links. These switches also have redundant uplink connections to the backbone switches, with each uplink port configured as a trusted routed uplink to provide faster convergence in case of an uplink failure.

The routers and Catalyst 6000 multilayer backbone switches have HSRP enabled for load balancing and redundant connectivity to guarantee mission-critical traffic.

Figure 1-3 Catalyst 3550 Switches in Wiring Closets in a Backbone Configuration

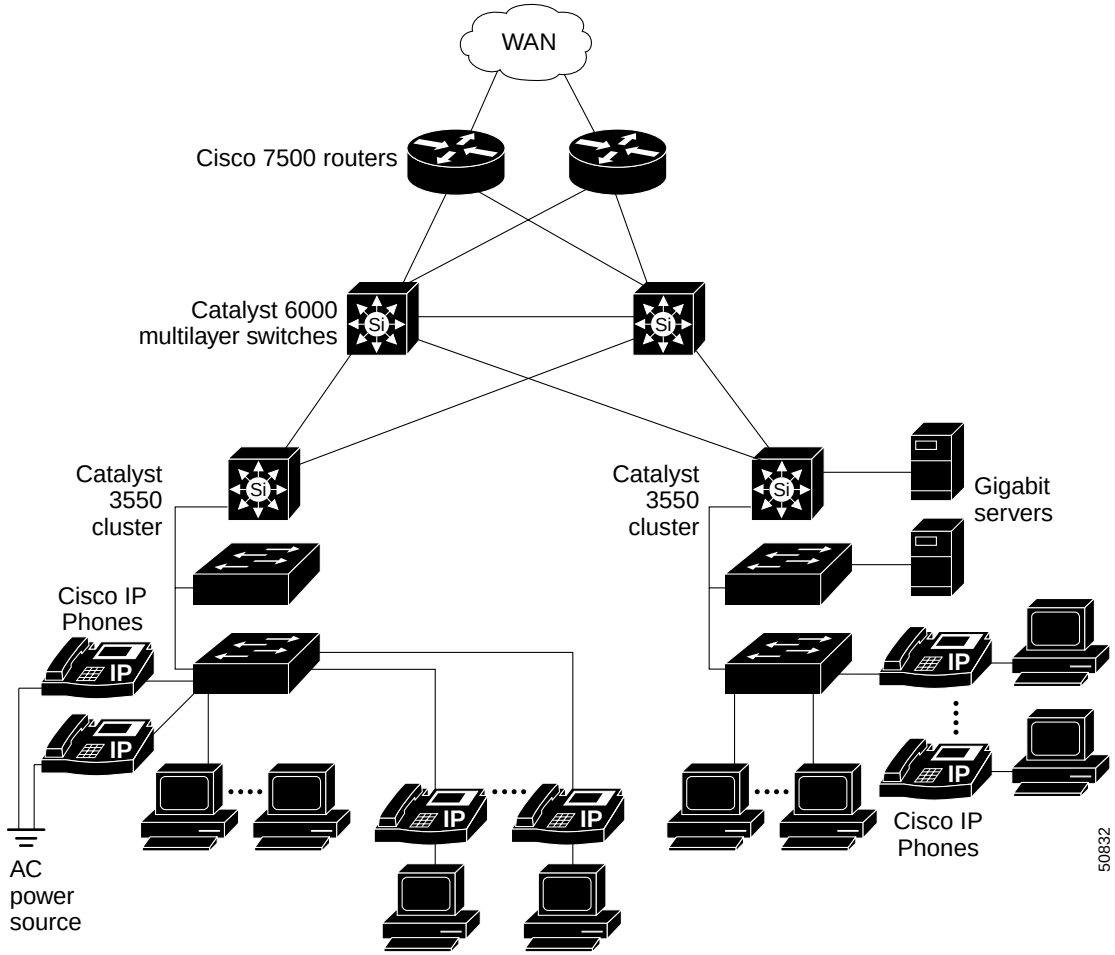

Multidwelling Network Using Catalyst 3550 Switches

A growing segment of residential and commercial customers are requiring high-speed access to Ethernet metropolitan-area networks (MANs). Figure 1-4 shows a configuration for a Gigabit Ethernet MAN ring using Catalyst 3550 multilayer switches as aggregation switches in the mini-point-of-presence (POP) location. These switches are connected through 1000BASE-X GBIC ports.

The resident switches can be Catalyst 3550 switches, providing customers with high-speed connections to the MAN. Catalyst 2912-LRE or 2924-LRE XL Layer 2-only switches also can be used as residential switches for customers requiring connectivity through existing phone lines. The Catalyst 2912-LRE or 2924-LRE XL switches can then connect to another residential switch or to an aggregation switch. For more information about the LRE switches, refer to the *Catalyst 2900 Series XL Hardware Installation Guide*.

All ports on the residential Catalyst 3550 switches (and Catalyst 2912-LRE XL or 2924-LRE XL switches if they are included) are configured as 802.1Q trunks with protected port and STP root guard features enabled. The protected port feature provides security and isolation between ports on the switch, ensuring that subscribers cannot view packets destined for other subscribers. STP root guard prevents

unauthorized devices from becoming the STP root switch. All ports have IGMP snooping or CGMP enabled for multicast traffic management. ACLs on the uplink ports to the aggregating Catalyst 3550 multilayer switches provide security and bandwidth management.

The aggregating switches and routers provide services such as those described in the previous examples, "Small to Medium-Sized Network Using Mixed Switches" and "Large Network Using Only Catalyst 3550 Switches."

Figure 1-4 Catalyst 3550 Switches in a MAN Configuration

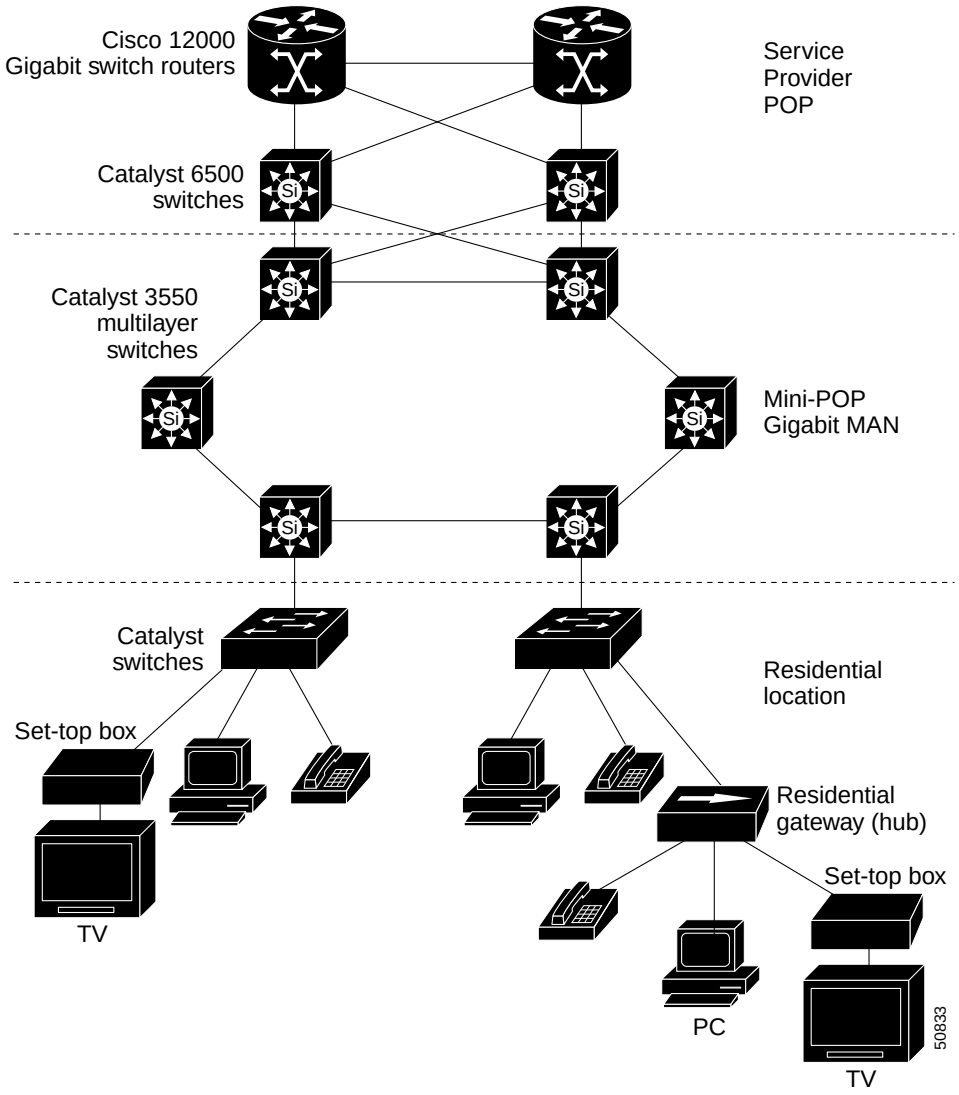

Using the Command-Line Interface

This chapter describes the IOS command-line interface (CLI) that you can use to configure your switches. It contains these sections:

- IOS Command Modes, page 2-1
- Getting Help, page 2-3
- Abbreviating Commands, page 2-3
- Using no and default Forms of Commands, page 2-4
- Understanding CLI Messages, page 2-4
- Using Command History, page 2-5
- Using Editing Features, page 2-6
- Searching and Filtering Output of show and more Commands, page 2-8
- Accessing the CLI, page 2-9

IOS Command Modes

The Cisco IOS user interface is divided into many different modes. The commands available to you depend on which mode you are currently in. Enter a question mark (?) at the system prompt to obtain a list of commands available for each command mode.

When you start a session on the switch, you begin in user mode, often called user EXEC mode. Only a limited subset of the commands are available in user EXEC mode. For example, most of the user EXEC commands are one-time commands, such as **show** commands, which show the current configuration status, and **clear** commands, which clear counters or interfaces. The user EXEC commands are not saved when the switch reboots.

To have access to all commands, you must enter privileged EXEC mode. Normally, you must enter a password to enter privileged EXEC mode. From this mode, you can enter any privileged EXEC command or enter global configuration mode.

Using the configuration modes (global, interface, and line), you can make changes to the running configuration. If you save the configuration, these commands are stored and used when the switch reboots. To access the various configuration modes, you must start at global configuration mode. From global configuration mode, you can enter interface configuration mode and line configuration mode.

Table 2-1 describes the main command modes, how to access each one, the prompt you see in that mode, and how to exit the mode. The examples in the table use the host name *Switch*.

Table 2-1 Command Mode Summary

Mode	Access Method	Prompt	Exit Method	About This Mode
User EXEC	Begin a session with your switch.	`Switch>`	Enter **logout** or **quit**.	Use this mode to • Change terminal settings. • Perform basic tests. • Display system information.
Privileged EXEC	While in user EXEC mode, enter the **enable** command.	`Switch#`	Enter **disable** to exit.	Use this mode to verify commands that you have entered. Use a password to protect access to this mode.
VLAN configuration	While in privileged EXEC mode, enter the **vlan database** command.	`Switch(vlan)#`	To exit to privileged EXEC mode, enter **exit**.	Use this mode to configure VLAN-specific parameters.
Global configuration	While in privileged EXEC mode, enter the **configure** command.	`Switch(config)#`	To exit to privileged EXEC mode, enter **exit** or **end**, or press **Ctrl-Z**.	Use this mode to configure parameters that apply to the entire switch.
Interface configuration	While in global configuration mode, enter the **interface** command (with a specific interface).	`Switch(config-if)#`	To exit to global configuration mode, enter **exit**. To return to privileged EXEC mode, press **Ctrl-Z** or enter **end**.	Use this mode to configure parameters for the Ethernet interfaces. To configure multiple interfaces with the same parameters, see the "Configuring a Range of Interfaces" section on page 8-9.
Line configuration	While in global configuration mode, specify a line with the **line vty** or **line console** command.	`Switch(config-line)#`	To exit to global configuration mode, enter **exit**. To return to privileged EXEC mode, press **Ctrl-Z** or enter **end**.	Use this mode to configure parameters for the terminal line.

Getting Help

You can enter a question mark (?) at the system prompt to display a list of commands available for each command mode. You can also obtain a list of associated keywords and arguments for any command, as shown in Table 2-2.

Table 2-2 Help Summary

Command	Purpose
help	Obtain a brief description of the help system in any command mode.

*abbreviated-command-entry***?**	Obtain a list of commands that begin with a particular character string.
	For example:
	`Switch# di?`
	`dir  disable  disconnect`

abbreviated-command-entry<**Tab**>	Complete a partial command name.
	For example:
	`Switch# sh conf<tab>`
	`Switch# show configuration`

?	List all commands available for a particular command mode. For example: `Switch> ?`
command **?**	List the associated keywords for a command. For example: `Switch> show ?`

command keyword **?**	List the associated arguments for a keyword.
	For example:
	`Switch(config)# cdp holdtime ?`
	`  <10-255>  Length of time (in sec) that receiver must keep this packet`

Abbreviating Commands

You have to enter only enough characters for the switch to recognize the command as unique. This example shows how to enter the **show configuration** privileged EXEC command:

```
Switch# show conf
```

Using no and default Forms of Commands

Almost every configuration command also has a **no** form. In general, use the **no** form to disable a feature or function or reverse the action of a command. For example, the **no shutdown** interface configuration command reverses the shutdown of an interface. Use the command without the keyword **no** to re-enable a disabled feature or to enable a feature that is disabled by default.

Configuration commands can also have a **default** form. The **default** form of a command returns the command setting to its default. Most commands are disabled by default, so the **default** form is the same as the **no** form. However, some commands are enabled by default and have variables set to certain default values. In these cases, the **default** command enables the command and sets variables to their default values.

Understanding CLI Messages

Table 2-3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 2-3 Common CLI Error Messages

Error Message	Meaning	How to Get Help
`% Ambiguous command: "show con"`	You did not enter enough characters for your switch to recognize the command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command are displayed.
`% Incomplete command.`	You did not enter all the keywords or values required by this command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command are displayed.
`% Invalid input detected at '^' marker.`	You entered the command incorrectly. The caret (^) marks the point of the error.	Enter a question mark (?) to display all the commands that are available in this command mode. The possible keywords that you can enter with the command are displayed.

Using Command History

The IOS provides a history or record of commands that you have entered. This feature is particularly useful for recalling long or complex commands or entries, including access lists. You can customize the command history feature to suit your needs as described in these sections:

- Changing the Command History Buffer Size, page 2-5
- Recalling Commands, page 2-5
- Disabling the Command History Feature, page 2-5

Changing the Command History Buffer Size

By default, the switch records ten command lines in its history buffer. Beginning in privileged EXEC mode, enter this command to change the number of command lines that the switch records during the current terminal session:

```
Switch# terminal history [size number-of-lines]
```

The range is from 0 to 256.

Beginning in line configuration mode, enter this command to configure the number of command lines the switch records for all sessions on a particular line:

```
Switch(config-line)# history [size number-of-lines]
```

The range is from 0 to 256.

Recalling Commands

To recall commands from the history buffer, perform one of the actions listed in Table 2-4:

Table 2-4 Recalling Commands

Action[1]	Result
Press **Ctrl-P** or the up arrow key.	Recall commands in the history buffer, beginning with the most recent command. Repeat the key sequence to recall successively older commands.
Press **Ctrl-N** or the down arrow key.	Return to more recent commands in the history buffer after recalling commands with **Ctrl-P** or the up arrow key. Repeat the key sequence to recall successively more recent commands.
show history	While in privileged EXEC mode, list the last several commands that you just entered. The number of commands that are displayed is determined by the setting of the **terminal history** global configuration command and **history** line configuration command.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Disabling the Command History Feature

The command history feature is automatically enabled.

To disable the feature during the current terminal session, enter the **terminal no history** privileged EXEC command.

To disable command history for the line, enter the **no history** line configuration command.

Using Editing Features

This section describes the editing features that can help you manipulate the command line. It contains these sections:

- Enabling and Disabling Editing Features, page 2-6
- Editing Commands through Keystrokes, page 2-6
- Editing Command Lines that Wrap, page 2-8

Enabling and Disabling Editing Features

Although enhanced editing mode is automatically enabled, you can disable it.

To re-enable the enhanced editing mode for the current terminal session, enter this command in privileged EXEC mode:

```
Switch# terminal editing
```

To reconfigure a specific line to have enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# editing
```

To globally disable enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# no editing
```

Editing Commands through Keystrokes

Table 2-5 shows the keystrokes that you need to edit command lines.

Table 2-5 Editing Commands through Keystrokes

Capability	Keystroke[1]	Purpose
Move around the command line to make changes or corrections.	Press **Ctrl-B**, or press the left arrow key.	Move the cursor back one character.
	Press **Ctrl-F**, or press the right arrow key.	Move the cursor forward one character.
	Press **Ctrl-A**.	Move the cursor to the beginning of the command line.
	Press **Ctrl-E**.	Move the cursor to the end of the command line.
	Press **Esc B**.	Move the cursor back one word.
	Press **Esc F**.	Move the cursor forward one word.
	Press **Ctrl-T**.	Transpose the character to the left of the cursor with the character located at the cursor.

Table 2-5 Editing Commands through Keystrokes (continued)

Capability	Keystroke[1]	Purpose
Recall commands from the buffer and paste them in the command line. The switch provides a buffer with the last ten items that you deleted.	Press **Ctrl-Y**.	Recall the most recent entry in the buffer.
	Press **Esc Y**.	Recall the next buffer entry. The buffer contains only the last 10 items that you have deleted or cut. If you press **Esc Y** more than ten times, you cycle to the first buffer entry.
Delete entries if you make a mistake or change your mind.	Press the **Delete** or **Backspace** key.	Erase the character to the left of the cursor.
	Press **Ctrl-D**.	Delete the character at the cursor.
	Press **Ctrl-K**.	Delete all characters from the cursor to the end of the command line.
	Press **Ctrl-U** or **Ctrl-X**.	Delete all characters from the cursor to the beginning of the command line.
	Press **Ctrl-W**.	Delete the word to the left of the cursor.
	Press **Esc D**.	Delete from the cursor to the end of the word.
Capitalize or lowercase words or capitalize a set of letters.	Press **Esc C**.	Capitalize at the cursor.
	Press **Esc L**.	Change the word at the cursor to lowercase.
	Press **Esc U**.	Capitalize letters from the cursor to the end of the word.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press **Ctrl-V** or **Esc Q**.	
Scroll down a line or screen on displays that are longer than the terminal screen can display. **Note** The More prompt is used for any output that has more lines than can be displayed on the terminal screen, including **show** command output. You can use the **Return** and **Space** bar keystrokes whenever you see the More prompt.	Press the **Return** key.	Scroll down one line.
	Press the **Space** bar.	Scroll down one screen.
Redisplay the current command line if the switch suddenly sends a message to your screen.	Press **Ctrl-L** or **Ctrl-R**.	Redisplay the current command line.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Editing Command Lines that Wrap

You can use a wraparound feature for commands that extend beyond a single line on the screen. When the cursor reaches the right margin, the command line shifts ten spaces to the left. You cannot see the first ten characters of the line, but you can scroll back and check the syntax at the beginning of the command.

To scroll back to the beginning of the command entry, press **Ctrl-B** or the left arrow key repeatedly. You can also press **Ctrl-A** to immediately move to the beginning of the line.

Note The arrow keys function only on ANSI-compatible terminals such as VT100s.

In this example, the **access-list** global configuration command entry extends beyond one line. When the cursor first reaches the end of the line, the line is shifted ten spaces to the left and redisplayed. The dollar sign ($) shows that the line has been scrolled to the left. Each time the cursor reaches the end of the line, the line is again shifted ten spaces to the left.

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1
Switch(config)# $ 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.25
Switch(config)# $t tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq
Switch(config)# $108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq 45
```

After you complete the entry, press **Ctrl-A** to check the complete syntax before pressing the **Return** key to execute the command. The dollar sign ($) appears at the end of the line to show that the line has been scrolled to the right:

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1$
```

The software assumes you have a terminal screen that is 80 columns wide. If you have a width other than that, use the **terminal width** privileged EXEC command to set the width of your terminal.

Use line wrapping with the command history feature to recall and modify previous complex command entries. For information about recalling previous command entries, see the "Editing Commands through Keystrokes" section on page 2-6.

Searching and Filtering Output of show and more Commands

You can search and filter the output for **show** and **more** commands. This is useful when you need to sort through large amounts of output or if you want to exclude output that you do not need to see.

To use this functionality, enter a **show** or **more** command followed by the *pipe* character (|), one of the keywords **begin**, **include**, or **exclude**, and an expression that you want to search for or filter out:

command **|** {**begin** | **include** | **exclude**} *regular-expression*

Expressions are case sensitive. For example, if you enter **| exclude output**, the lines that contain *output* are not displayed, but the lines that contain *Output* are displayed.

This example shows how to include in the output display only lines where the expression *protocol* appears:

```
Switch# show interfaces | include protocol
Vlan1 is up, line protocol is up
Vlan10 is up, line protocol is down
GigabitEthernet0/1 is up, line protocol is down
GigabitEthernet0/2 is up, line protocol is up
```

Accessing the CLI

Before you can access the CLI, you need to connect a terminal or PC to the switch console port and power on the switch as described in the hardware installation guide that shipped with your switch. Then, to understand the boot process and the options available for assigning IP information, see Chapter 4, "Assigning the Switch IP Address and Default Gateway."

If your switch is already configured, you can access the CLI through a local console connection or through a remote Telnet session, but your switch must first be configured for this type of access. For more information, see the "Setting a Telnet Password for a Terminal Line" section on page 6-6.

Getting Started with CMS

This chapter provides these topics about the Cluster Management Suite (CMS) software:

- Features, page 3-2
- Front Panel View, page 3-4
- Topology View, page 3-10
- Menus and Toolbar, page 3-15
- Interaction Modes, page 3-26
- Wizards, page 3-26
- Online Help, page 3-27
- CMS Window Components, page 3-28
- Accessing CMS, page 3-30
- Verifying Your Changes, page 3-32
- Saving Your Changes, page 3-32
- Using Different Versions of CMS, page 3-33
- Where to Go Next, page 3-33

Note

- For system requirements and for browser and Java plug-in configuration procedures, refer to the release notes.
- For procedures for using CMS, refer to the online help.

Note This chapter describes the CMS interface of the Catalyst 3550 switches. Refer to the appropriate switch documentation for descriptions of the web-based management software used on other Catalyst switches.

Features

CMS provides these features (Figure 3-1) for managing switch clusters and individual switches from Web browsers such as Netscape Communicator or Microsoft Internet Explorer:

- Two views of your network that can be displayed at the same time:
  - The Front Panel view displays the front-panel image of a specific switch or the front-panel images of all switches in a cluster. From this view, you can select multiple ports or multiple switches and configure them with the same settings.

    When CMS is launched from a command switch, the Front Panel view displays the front-panel images of all switches in the cluster. When CMS is launched from a noncommand switch, the Front Panel view displays only the front panel of the specific switch.

Note CMS from a standalone switch or from a noncommand switch is referred to as *Device Manager* (also referred to as *Switch Manager*). Device Manager is for configuring an individual switch. When you select Device Manager for a specific switch in the cluster, you launch a separate CMS session. The Device Manager interface can vary between the Catalyst switch platforms.

  - The Topology view displays a network map that uses icons that represent switch clusters, cluster members, cluster candidates, neighboring devices that are not eligible to join a cluster, and link types. From this view, you can select multiple switches and configure them to run with the same settings. You can also display link information in the form of link reports and link graphs.

    This view is available only when CMS is launched from a command switch.
- Menus and toolbar to access configuration and management options:
  - The menu bar provides the complete list of options for managing a single switch and switch clusters.
  - The toolbar provides buttons for commonly used switch and cluster configuration options and information windows such as legends and online help.
  - The port popup menu, in the Front Panel view, provides options specific for configuring and monitoring switch ports.
  - The device popup menu, in either the Front Panel or the Topology views, provides switch and cluster configuration and monitoring options.
  - The candidate, member, and link popup menus provide options for configuring and monitoring devices and links in the Topology view.

  The toolbar and popup menus provide quick ways to access frequently used menu-bar options.
- Tools to simplify configuration tasks:
  - Interactive modes—guide mode and expert mode—that control the presentation of some complex configuration options
  - Wizards that require minimal information from you to configure some complex features
  - Comprehensive online help that provides high-level concepts and procedures for performing tasks from the window

- Two levels of access to the configuration options: read-write access for users allowed to change switch settings; read-only access for users allowed to only view switch settings
- Consistent set of GUI components (such as tabs, buttons, drop-down lists, tables, and so on) for a consistent approach to setting configuration parameters

Figure 3-1 CMS Features

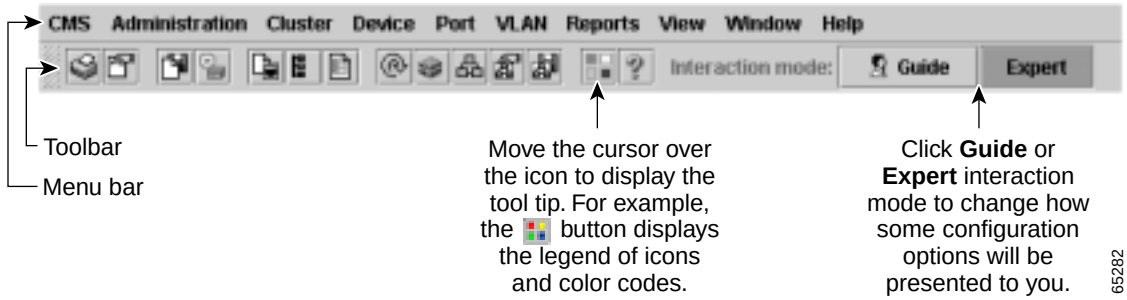

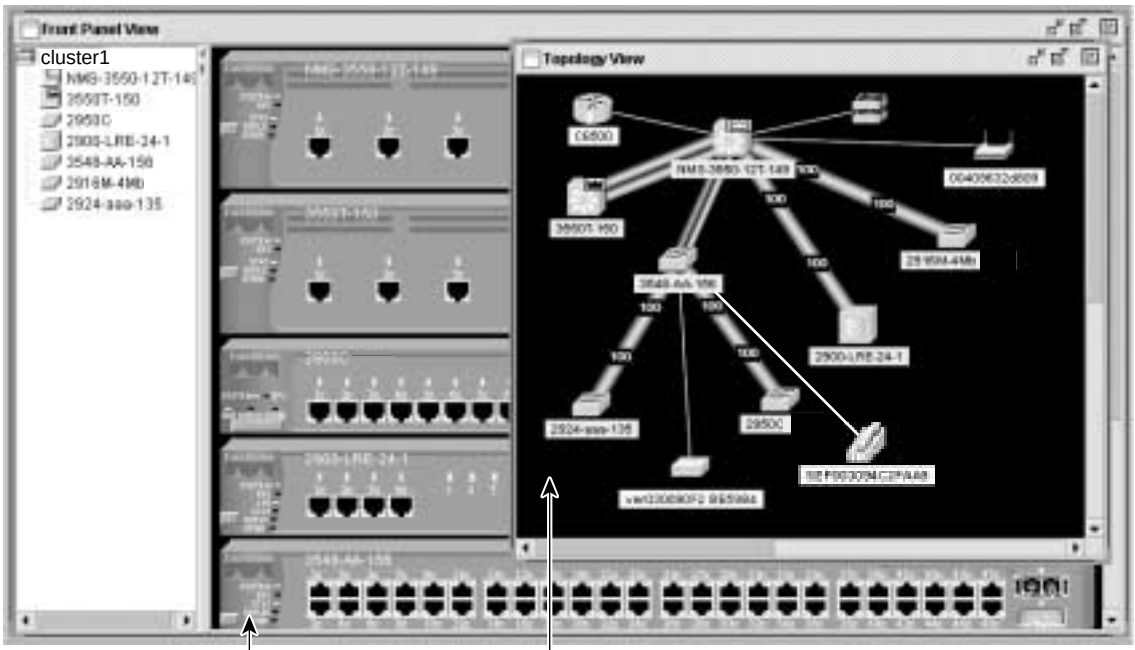

Front Panel View

When CMS is launched from a command switch, the Front Panel view displays the front-panel images of all switches in the cluster (Figure 3-2). When CMS is launched from a standalone or non-command member switch, the Front Panel view displays only the front panel of the specific switch (Figure 3-3).

Figure 3-2 Front Panel View from a Standalone Switch

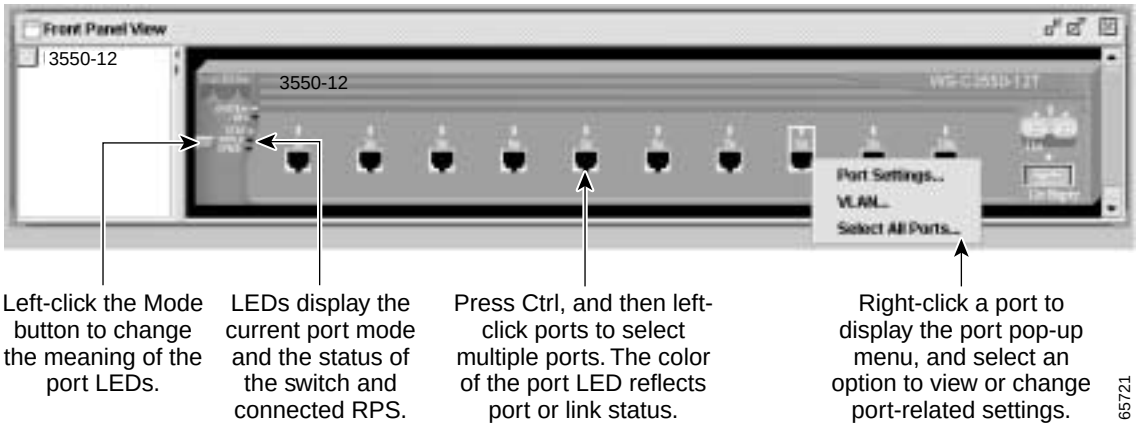
Front Panel View
3550-12
3550-12
Port Settings...
VLAN...
Select All Ports...
Left-click the Mode button to change the meaning of the port LEDs.
LEDs display the current port mode and the status of the switch and connected RPS.
Press Ctrl, and then left-click ports to select multiple ports. The color of the port LED reflects port or link status.
Right-click a port to display the port pop-up menu, and select an option to view or change port-related settings.
65721

Figure 3-3 Front Panel View from a Command Switch

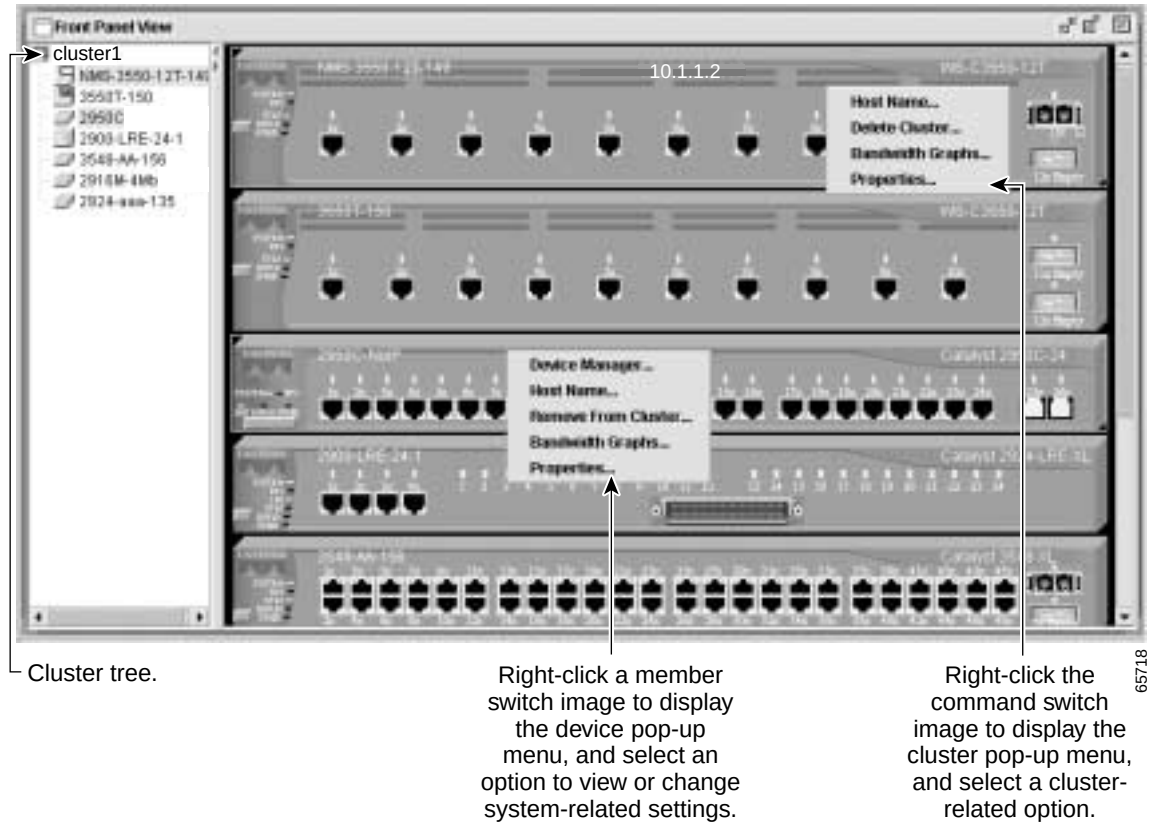

Cluster Tree

The cluster tree (Figure 3-3) appears in the left frame of the Front Panel view and shows the name of the cluster and a list of its members. The sequence of the cluster-tree icons (Figure 3-4) mirror the sequence of the front-panel images. You can change the sequence by selecting **View > Arrange Front Panel**. The colors of the devices in the cluster tree reflect the status of the devices (Table 3-1).

If you want to configure switch or cluster settings on one or more switches, select the appropriate front-panel images.

- To select a front-panel image, click either the cluster-tree icon or the corresponding front-panel image. The front-panel image is then highlighted with a yellow outline.
- To select multiple front-panel images, press the **Ctrl** key, and left-click the cluster-tree icons or the front-panel images. To deselect an icon or image, press the **Ctrl** key, and left-click the icon or image.

If the cluster has many switches, you might need to scroll down the window to display the rest of front-panel images. Instead of scrolling, you can click an icon in the cluster tree, and CMS then scrolls and displays the corresponding front-panel image.

Figure 3-4 Cluster-Tree Icons

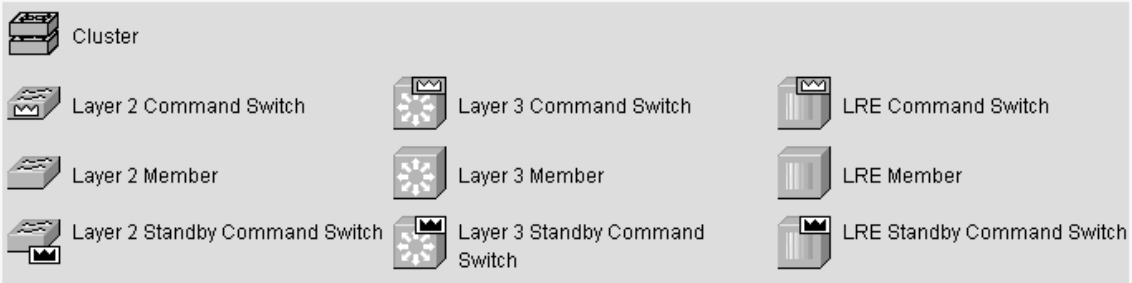

Table 3-1 Cluster Tree Icon Colors

Color	Device Status
Green	Switch is operating normally.
Yellow	The internal fan of the switch is not operating, or the switch is receiving power from an RPS.
Red	Switch is not powered up, has lost power, or the command switch is unable to communicate with the member switch.

Front-Panel Images

You can manage the switch from a remote station by using the front-panel images. The front-panel images are updated based on the network polling interval that you set from **CMS > Preferences**.

Note The Preferences window is not available if your switch access level is read-only. For more information about the read-only access mode, see the "Access Modes in CMS" section on page 3-31.

Figure 3-5 shows the port icons as they appear in the front-panel images. To select a port, click the port on the front-panel image. The port is then highlighted with a yellow outline. To select multiple ports, you can:

- Press the left mouse button, drag the pointer over the group of ports that you want to select, and then release the mouse button.
- Press the **Ctrl** key, and click the ports that you want to select.
- Right-click a port, and select **Select All Ports** from the port popup menu.

Figure 3-5 Port Icons

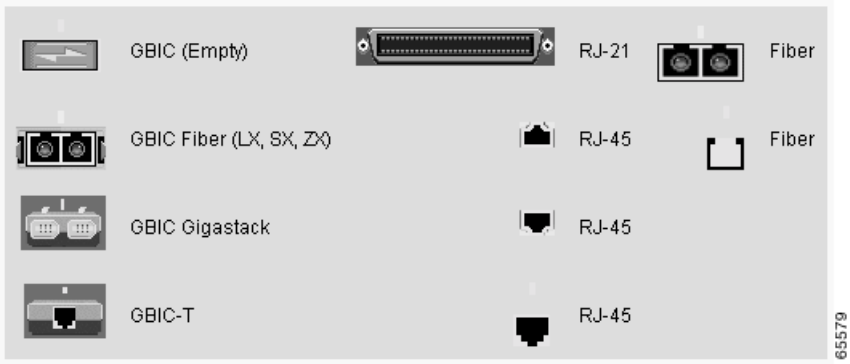

The following sections provide complete descriptions of the LED images. Similar descriptions of these LEDs are provided in the switch hardware installation guide.

Redundant Power System LED

The Redundant Power System (RPS) LED shows the RPS status (Table 3-2). Certain switches in the switch cluster use a specific RPS model:

- Cisco RPS 300 (model PWR300-AC-RPS-N1)—Catalyst 2900 LRE XL, Catalyst 2950, Catalyst 3524-PWR XL, and Catalyst 3550 switches
- Cisco RPS 600 (model PWR600-AC-RPS)—Catalyst 2900 XL and Catalyst 3500 XL switches, except the Catalyst 2900 LRE XL and Catalyst 3524-PWR XL switches

Refer to the appropriate switch hardware documentation for RPS descriptions specific for the switch.

Table 3-2 RPS LED

Color	RPS Status
Black (off)	RPS is off or is not installed.
Green	RPS is connected and operational.
Blinking green	RPS is providing power to another switch in the stack.
Amber	RPS is connected but not functioning. The RPS could be in standby mode. To put the RPS in Active mode, press the Standby/Active button on the RPS, and the LED should turn green. If it does not, one of these conditions could exist: • One of the RPS power supplies could be down. Contact Cisco Systems. • The RPS fan could have failed. Contact Cisco Systems.
Blinking amber	Internal power supply of the switch is down, and redundancy is lost. The switch is operating on the RPS.

Port Modes and LEDs

The port modes (Table 3-3) determine the type of information displayed through the port LEDs. When you change port modes, the meanings of the port LED colors (Table 3-4) also change.

Note The bandwidth utilization mode (UTL LED) does not appear on the front-panel images. Select **Reports > Bandwidth Graphs** to display the total bandwidth in use by the switch. Refer to the switch hardware installation guide for information about using the UTL LED.

To select or change a mode, click the Mode button until the desired mode LED is green.

Table 3-3 Port Modes

Mode LED	Description
STAT	Link status of the ports. Default mode.
DUPLX	Duplex setting on the ports.
SPEED	Speed setting on the ports.

Table 3-4 Port LEDs

Port Mode	Port LED Color	Description
STAT	Cyan (off)	No link.
	Green	Link present.
	Amber	Link fault. Error frames can affect connectivity, and errors such as excessive collisions, CRC errors, and alignment and jabber errors are monitored for a link-fault indication. Port is not forwarding. Port was disabled by management, by an address violation, or was blocked by Spanning Tree Protocol (STP). **Note** After a port is reconfigured, the port LED can remain amber for up to 30 seconds as STP checks the switch for possible loops.
	Brown	No link and port is administratively shut down.
DUPLX	Cyan (off)	Port is operating in half-duplex mode.
	Green	Port is operating in full-duplex mode.
SPEED	Cyan (off)	Port is operating at 10 Mbps or no link.
	Green	Port is operating at 100 Mbps (10/100 ports), 155 Mbps (ATM ports), or 1000 Mbps (fixed Gigabit port).
	Blinking green	Port is operating at 1000 Mbps (10/100/1000 ports).

VLAN Membership Modes

Ports in the Front Panel view are outlined by colors (Table 3-5) when you click **Highlight VLAN Port Membership Modes** on the Configure VLANs tab on the VLAN window (**VLAN > VLAN > Configure VLANs**). The colors show the VLAN membership mode of each port. The VLAN membership mode determines the kind of traffic the port carries and the number of VLANs it can belong to. For more information about these modes, see the "VLAN Port Membership Modes" section on page 9-3.

Note This feature is not supported on the Catalyst 1900 and Catalyst 2820 switches.

Table 3-5 VLAN Membership Modes

Mode	Color
Static access	Light green
Dynamic access	Pink
ISL trunk	Orange
802.1Q trunk	Peach
Negotiate trunk	White

Topology View

The Topology view displays how the devices within a switch cluster are connected and how the switch cluster is connected to other clusters and devices. From this view, you can add and remove cluster members. This view provides two levels of detail of the network topology:

- When you right-click a cluster icon and select Expand Cluster, the Topology view displays the switch cluster in detail. This view shows the command switch and member switches in a cluster. It also shows candidate switches that can join the cluster. This view does not display the details of any neighboring switch clusters (Figure 3-6).
- When you right-click a command-switch icon and select Collapse Cluster, the cluster is collapsed and represented by a single icon. The view shows how the cluster is connected to other clusters, candidate switches, and devices that are not eligible to join the cluster (such as routers, access points, IP phones, and so on) (Figure 3-7).

Note The Topology view displays only the switch cluster and network neighborhood of the specific command or member switch that you access. To display a different switch cluster, you need to access the command switch or member switch of that cluster.

You can arrange the device icons in this view. To move a device icon, click and drag the icon. To select multiple device icons, you can either:

- Press the left mouse button, drag the pointer over the group of device icons that you want to select, and then release the mouse button.
- Press the **Ctrl** key, and click the device icons that you want to select.

After selecting the icons, drag the icons to any area in the view.

Figure 3-6 Expand Cluster View

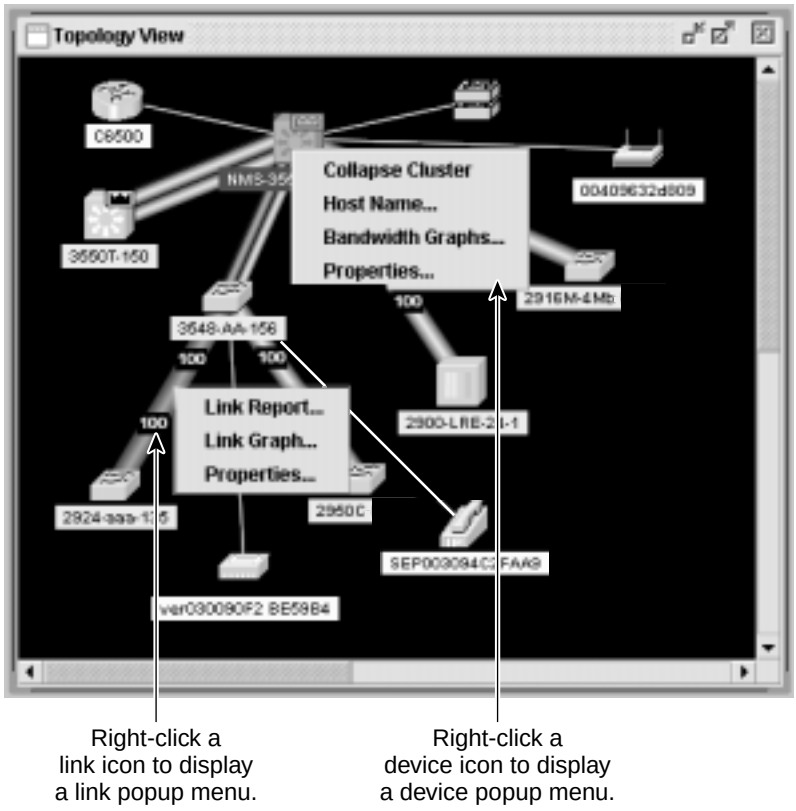

Figure 3-7 Collapse Cluster View

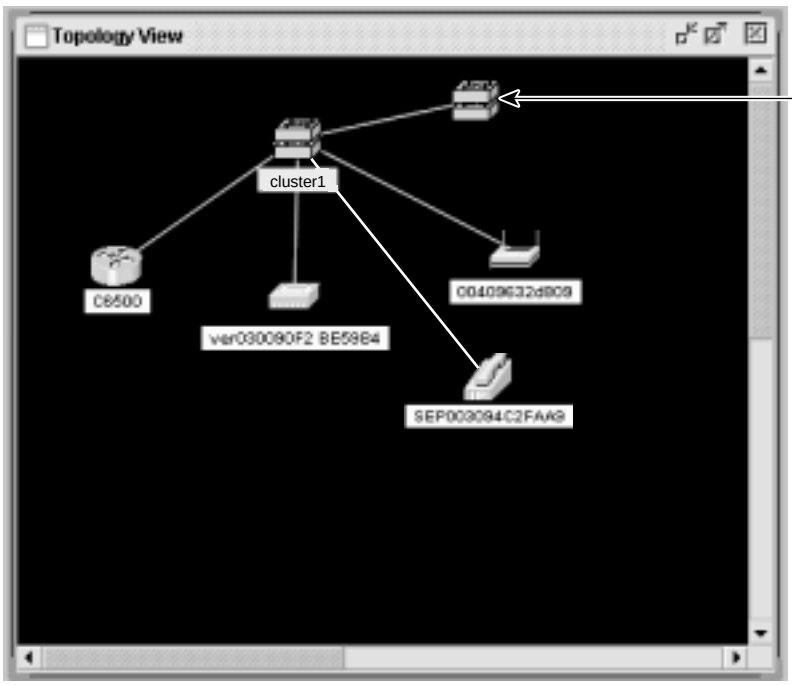

Topology Icons

The Topology view and the cluster tree use the same set of device icons to represent clusters, command and standby command switches, and member switches (Figure 3-8). The Topology view also uses additional icons to represent these types of neighboring devices:

- Customer premises equipment (CPE) devices that are connected to Long-Reach Ethernet (LRE) switches
- Devices that are not eligible to join the cluster, such as Cisco IP phones, Cisco access points, and CDP-capable hubs and routers
- Devices that are identified as *unknown* devices, such as some Cisco devices and third-party devices

Note Candidate switches are distinguished by the color of their device label. Device labels and their colors are described in the "Colors in the Topology View" section on page 3-14.

To select a device, click the icon. The icon is then highlighted. To select multiple devices, you can either:

- Press the left mouse button, drag the pointer over the group of icons that you want to select, and then release the mouse button.
- Press the **Ctrl** key, and click the icons that you want to select.

Figure 3-8 Topology-View Device Icons

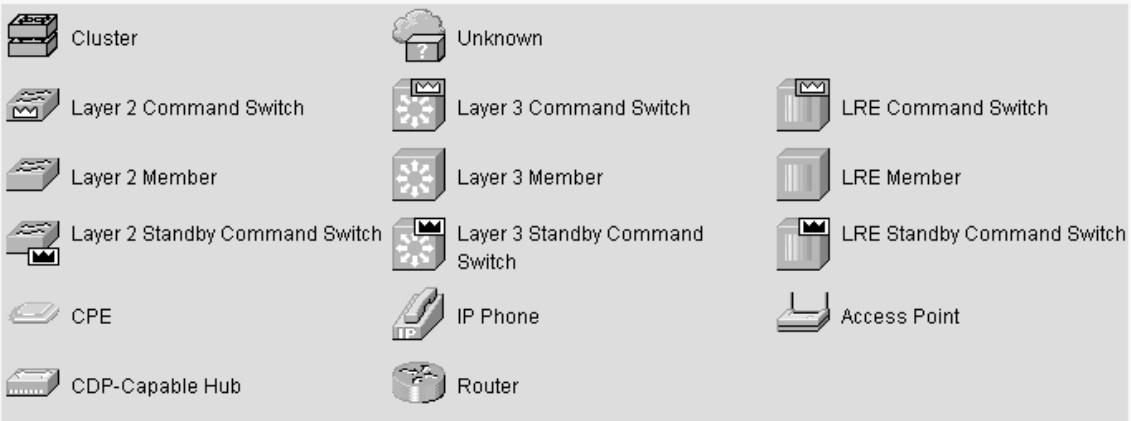

The Topology view also uses a set of link icons (Figure 3-9) to show the link type and status between two devices. To select a link, click the link that you want to select. To select multiple links, press the Ctrl key, and click the links that you want to select.

Figure 3-9 Topology-View Link Icons

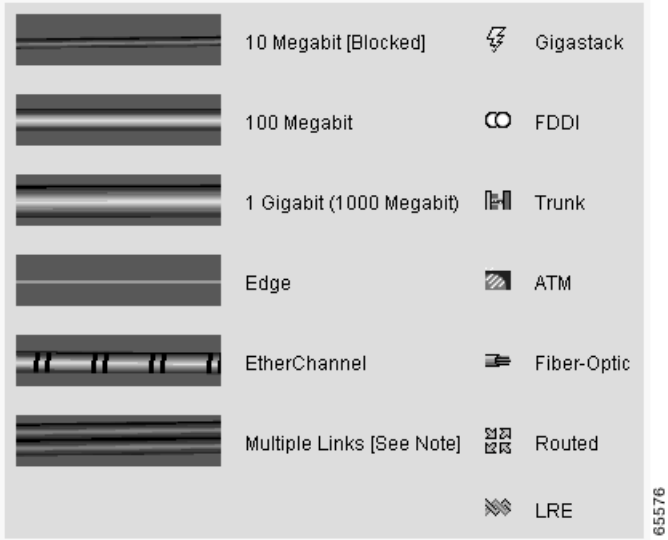

Device and Link Labels

The Topology view displays device and link information by using these *labels*:

- Cluster and switch names
- Switch MAC and IP addresses
- Link type between the devices
- Link speed and IDs of the interfaces on both ends of the link

When using these labels, keep these considerations in mind:

- The IP address displays only in the labels for the command switch and member switches.
- The label of a neighboring cluster icon only displays the IP address of the command-switch IP address.
- The displayed link speeds are the actual link speeds except on the LRE links, which display the administratively assigned speed settings.

You can change the label settings from the Topology Options window, which is displayed by selecting **View > Topology Options**.

Colors in the Topology View

The colors of the Topology view icons reflect the status of the devices and links (Table 3-6, Table 3-7, and Table 3-8).

Table 3-6 Device Icon Colors

Icon Color	Color Meaning
Green	The device is operating.
Yellow[1]	The internal fan of the switch is not operating, or the switch is receiving power from an RPS.
Red[1]	The device is not operating.

1. Available only on the cluster members.

Table 3-7 Single Link Icon Colors

Link Color	Color Meaning
Green	Active link
Red	Down or blocked link

Table 3-8 Multiple Link Icon Colors

Link Color	Color Meaning
Both green	All links are active.
One green; one red	One link is active, and at least one link is down or blocked.
Both red	All links are down or blocked.

The color of a device label shows the cluster membership of the device (Table 3-9).

Table 3-9 Device Label Colors

Label Color	Color Meaning
Green	A cluster member, either a member switch or the command switch
Cyan	A candidate switch that is eligible to join the cluster
Yellow	An unknown device or a device that is not eligible to join the cluster

Topology Display Options

You can set the type of information displayed in the Topology view by changing the settings in the Topology Options window. To display this window, select **View > Topology Options**. From this window, you can select:

- Device icons to be displayed in the Topology view
- Labels to be displayed with the device and link icons

Menus and Toolbar

The configuration and monitoring options for configuring switches and switch clusters are available from the menu bar, toolbar, and the Front-Panel and Topology view popup menus.

Menu Bar

The menu bar provides the complete list of options for managing a single switch and switch cluster. The menu bar is the same whether or not the Front-Panel or Topology views are displayed.

Options displayed from the menu bar can vary:

- The option for enabling a command switch is only available from a CMS session launched from a command-capable switch.
- Cluster management tasks, such as upgrading the software of groups of switches, are available only from a CMS session launched from a command switch.
- If you launch CMS from a specific switch, the menu bar displays the features supported only by that switch.
- If you launch CMS from a command switch, the menu bar displays the features supported on the switches in the cluster, with these exceptions:
  - If the command switch is a Layer 3 switch, such as a Catalyst 3550 switch, the menu bar displays the features of all Layer 3 and Layer 2 switches in the cluster.
  - If the command switch is a Layer 2 switch, such as a Catalyst 2950 or Catalyst 3500 XL switch, the menu bar displays the features of all Layer 2 switches in the cluster. The menu bar does not display Layer 3 features even if the cluster has Catalyst 3550 Layer 3 member switches.

Note We strongly recommend that the highest-end, command-capable switch in the cluster be the command switch so that all of the features supported in the cluster are displayed from the menu bar. If you have a switch cluster with a Catalyst 3550, that switch should be the command switch. If your switch cluster has Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches, the Catalyst 2950 should be the command switch. Refer to the release notes for the Catalyst switches that can be part of a switch cluster.

Note

- Unless noted otherwise, Table 3-10 lists the menu-bar options available from a Catalyst 3550 command switch and when the cluster contains *only* Catalyst 3550 member switches. The menu bar of the command switch displays all menu-bar options available from the cluster, including options from member switches from other cluster-capable switch platforms.
- The menu-bar options on a Catalyst 3550 switch change depending on whether the switch is running the enhanced multilayer software image or not. Footnotes describe the availability of an option if the switch is running the enhanced multilayer software image.
- The footnotes in the table describe the availability of an option based on your access mode in CMS: read-only (access level 1–14) and read-write (access level 15). For more information about CMS access modes, see the "Access Modes in CMS" section on page 3-31.

- If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:
  - Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 3550 member switches running Cisco IOS Release 12.1(6)EA1 or earlier

  For more information about this limitation, refer to the Catalyst 3550 release notes.
- These switches do not support CMS access modes:
  - Catalyst 1900 and Catalyst 2820
  - Catalyst 2900 XL switches with 4-MB CPU DRAM

  If these switches are in a cluster that is using CMS access levels, these switches appear as unavailable devices and cannot be configured from CMS.

Table 3-10 Menu Bar

Menu-Bar Options	Task
CMS	
Page Setup	Set default document printer properties to be used when printing from CMS.
Print Preview	View the way the CMS window or help file will appear when printed.
Print	Print a CMS window or help file.
Guide Mode/Expert Mode[1]	Select which interaction mode to use when you select a configuration option.
Preferences[2]	Set CMS display properties, such as polling intervals, the default views to open at startup, and the color of administratively shutdown ports.
Administration	
IP Addresses[2]	Configure IP information for a switch.
SNMP[2]	Enable and disable Simple Network Management Protocol (SNMP), enter community strings, and configure end stations as trap managers.
System Time[2]	Configure the system time or configure the Network Time Protocol (NTP).
HTTP Port[2]	Configure the Hypertext Transfer Protocol (HTTP) port.
Console Baud Rate[2]	Change the baud rate for the switch console port.
MAC Addresses[2]	Enter dynamic, secure, and static addresses in a switch address table. You can also define the forwarding behavior of static addresses.
ARP[2]	Display the device Address Resolution Protocol (ARP) table, and configure the ARP cache timeout setting.
Save Configuration[1]	Save the configuration for the cluster or switch to Flash memory.
Software Upgrade[1]	Upgrade the software for the cluster or a switch.
System Reload[1]	Reboot the switch with the latest installed software.

Table 3-10 Menu Bar (continued)

Menu-Bar Options	Task
Cluster	
Cluster Manager[3]	Launch a CMS session from the command switch.
Create Cluster[1 4]	Designate a command switch, and name a cluster.
Delete Cluster[1 5]	Delete a cluster.
Add to Cluster[1 5]	Add a candidate to a cluster.
Remove from Cluster[1 5]	Remove a member from the cluster.
Standby Command Switches[2 5]	Create a Hot Standby Router Protocol (HSRP) standby group to provide command-switch redundancy.
Hop Count[2 5]	Enter the number of hops away that a command switch looks for members and for candidate switches.
Device	
Device Manager[5]	Launch Device Manager for a specific switch.
Host Name[1]	Change the host name of a switch.
STP[2]	Display and configure STP parameters for a switch.
IGMP Snooping[2]	Enable and disable Internet Group Management Protocol (IGMP) snooping and IGMP Immediate-Leave processing on the switch. Join or leave multicast groups, and configure multicast routers.
ACL[2] (guide mode available[1])	Create and maintain access control lists (ACLs), and attach ACLs to specific ports.
Security Wizard[1]	Filter certain traffic, such as HTTP traffic, to certain users or devices.
QoS[2] (guide mode available on some options[1])	Display submenu options to enable and disable quality of service (QoS) and to configure or modify these parameters: • Trust enable/disable[2] • Trust settings[2] • Queues[2] • Maps[2] • Classes[2] (guide mode available[1]) • Aggregate policers[2] (guide mode available[1]) • Policies[2] (guide mode available[1]) • Statistics[2]
IP Routing[2 6] (guide mode available[1])	Display submenu options to configure or modify these parameters: • Protocols[2] (guide mode available[1]) • Static routing[2] (guide mode available[1]) • Equal- and unequal-cost routing[2] (guide mode available[1])
IP Multicast Wizard[1 6]	Provide minimum information to configure IP multicast routing on a device so that it can forward multicast packets as a part of a multicast tree.
IP Multicast Routing[2 6]	Enable and configure multicast routing.

Table 3-10 Menu Bar (continued)

Menu-Bar Options	Task
Router Redundancy[2 6] (guide mode available[1])	Add a switch to or remove a switch from an HSRP group.
Fallback Bridging[2 6]	Create a fallback bridging group, modify a group, delete a group, or view its details.
802.1X[1]	Configure 802.1X authentication of devices as they are attached to LAN ports in a point-to-point infrastructure.
AVVID Wizards[1]	• Video Wizard[1]—Optimize multiple video servers for transmitting video traffic. • Priority Data Wizard[1]—Provide a higher priority to specific applications.
Port	
Port Settings[2]	Display and configure port parameters on a switch.
Port Search	Search for a port through its description.
Port Security[1]	Enable port security on a port.
EtherChannels[2]	Group ports into logical units for high-speed links between switches.
SPAN[2]	Enable Switch Port Analyzer (SPAN) port monitoring.
Protected Port[2]	Configure a port to prevent it from receiving bridged traffic from another port on the same switch.
Flooding Control[2]	Block the normal flooding of unicast and multicast packets, and enable the switch to block packet storms.
VLAN	
VLAN[2] (guide mode available[1])	Display VLAN membership, assign ports to VLANs, and configure Inter-Switch Link (ISL) and 802.1Q trunks. Display and configure the VLAN Trunking Protocol (VTP) for interswitch VLAN membership.
VMPS[2]	Configure the VLAN Membership Policy Server (VMPS).
VLAN Maps[2]	Configure VLAN maps.

Table 3-10 Menu Bar (continued)

Menu-Bar Options	Task
Reports	
Inventory	Display the device type, software version, IP address, and other information about a switch.
Port Statistics	Display port statistics.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use by the switch.
Link Graphs	Display a graph showing the bandwidth being used for the selected link.
Link Reports	Display the link report for two connected devices. If one device is an unknown device or a candidate, only the cluster-member side of the link displays.
QoS Reports	Display QoS reports of incoming or outgoing traffic for specific device interfaces.
QoS Graphs	Display QoS graphs of incoming or outgoing traffic for specific device interfaces.
Router Reports[6]	Display reports with an excerpt from the routing table on the switch and the attributes of the HSRP group in which the switch participates.
Fallback Bridging[6]	Display a report of all fallback bridging groups and their attributes.
System Messages	Display the most recent system messages (IOS messages and switch-specific messages) sent by the switch software. This option is available on the Catalyst 2950 or Catalyst 3550 switches. It is not available from the Catalyst 2900 XL and Catalyst 3500 XL switches. You can display the system messages of the Catalyst 2900 XL and Catalyst 3500 XL switches when they are in a cluster where the command switch is a Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later or a Catalyst 3550 switch running Cisco IOS Release 12.1(8)EA1 or later. For more information about system messages, refer to the switch system messages guide.
View	
Refresh	Update the views with the latest status.
Front Panel	Display the Front Panel view.
Arrange Front Panel[1 5]	Rearrange the order in which switches appear in the Front Panel view.
Topology[5]	Display the Topology view.
Topology Options [5]	Select the information to be displayed in the Topology view.
Automatic Topology Layout[5]	Request CMS to rearrange the topology layout.
Save Topology Layout[1 5]	Save the presentation of the cluster icons that you arranged in the Topology view to Flash memory.
Window	List the open windows in your CMS session.

Table 3-10 Menu Bar (continued)

Menu-Bar Options	Task
Help	
Overview	Obtain an overview of the CMS interface.
What's New	Obtain a description of the new CMS features.
Help For Active Window	Display the help for the active open window. This is the same as clicking **Help** from the active window.
Contents	List all of the available online help topics.
Legend	Display the legend that describes the icons, labels, and links.
About	Display the CMS version number.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.
2. Some options from this menu option are not available in read-only mode.
3. Available only from a Device Manager session on a cluster member.
4. Available only from a Device Manager session on a command-capable switch that is not a cluster member.
5. Available only from a cluster management session.
6. Available only from a switch running the enhanced multilayer software image.

Toolbar

The toolbar buttons display commonly used switch and cluster configuration options and information windows such as legends and online help. Hover the cursor over an icon to display the feature. Table 3-11 describes the toolbar options, from left to right on the toolbar.

Table 3-11 Toolbar Buttons

Toolbar Option	Keyboard Shortcut	Task
Print	Ctrl-P	Print a CMS window or help file.
Preferences[1]	Ctrl-R	Set CMS display properties, such as polling intervals, the views to open at CMS startup, and the color of administratively shutdown ports.
Save Configuration[2]	Ctrl-S	Save the configuration for the cluster or switch to Flash memory.
Software Upgrade[2]	Ctrl-U	Upgrade the software for the cluster or a switch.
Port Settings[1]	–	Display and configure port parameters on a switch.
VLAN[1]	–	Display VLAN membership, assign ports to VLANs, and configure ISL and 802.1Q trunks.
Inventory	–	Display the device type, the software version, the IP address, and other information about a switch.
Refresh	–	Update the views with the latest status.
Front Panel	–	Display the Front Panel view.
Topology[3]	–	Display the Topology view.
Topology Options[3]	–	Select the information to be displayed in the Topology view.
Save Topology Layout[2] [3]	–	Save the presentation of the cluster icons that you arranged in the Topology view to Flash memory.
Legend	–	Display the legend that describes the icons, labels, and links.
Help For Active Window	F1 key	Display the help for the active open window. This is the same as clicking **Help** from the active window.

1. Some options from this menu option are not available in read-only mode.
2. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.
3. Available only from a cluster-management session.

Front Panel View Popup Menus

These popup menus are available in the Front Panel view.

Device Popup Menu

You can display all switch and cluster configuration windows from the menu bar, or you can display commonly used configuration windows from the device popup menu (Table 3-12). To display the device popup menu, click the switch icon from the cluster tree or the front-panel image itself, and right-click.

Table 3-12 Device Popup Menu

Popup Menu Option	Task
Device Manager[1]	Launch Device Manager for the switch.
Delete Cluster[2 3 4]	Delete a cluster.
Remove from Cluster[3 4]	Remove a member from the cluster.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use.
Host Name[4]	Change the name of the switch.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Available from a cluster member switch but not from the command switch.
2. Available only from the command switch.
3. Available only from a cluster-management session.
4. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.

Port Popup Menu

You can display all port configuration windows from the **Port** menu on the menu bar, or you can display commonly used port configuration windows from the port popup menu (Table 3-13). To display the port popup menu, click a specific port image, and right-click.

Table 3-13 Port Popup Menu

Popup Menu Option	Task
Port Settings[1]	Display and configure port settings.
VLAN[1]	Define the VLAN mode for a port or ports and add ports to VLANs. Not available for the Catalyst 1900 and Catalyst 2820 switches.
Port Security[1 2]	Enable port security on a port.
Link Graphs[3]	Display a graph showing the bandwidth used by the selected link.
Select All Ports	Select all ports on the switch for global configuration.

1. Some options from this menu option are not available in read-only mode.
2. Available on switches that support the Port Security feature.
3. Available only when there is an active link on the port (that is, the port LED is green when in port status mode).

Topology View Popup Menus

These popup menus are available in the Topology view.

Link Popup Menu

You can display reports and graphs for a specific link displayed in the Topology view (Table 3-14). To display the link popup menu, click the link icon, and right-click.

Table 3-14 Link Popup Menu

Popup Menu Option	Task
Link Report	Display the link report for two connected devices. If one device is an unknown device or a candidate, only the cluster member side of the link displays.
Link Graph	Display a graph showing the bandwidth used by the selected link.
Properties	Display information about the device and port on either end of the link and the state of the link.

The Link Report and Link Graph options are not available if at both ends of the link are

- Candidate switches
- Catalyst 1900 and Catalyst 2820 switches
- Devices that are not eligible to join the cluster

If multiple links are configured between two devices, when you click the link icon and right-click, the Multilink Content window appears (Figure 3-10). Click the link icon in this window, and right-click to display the link popup menu specific for that link.

Figure 3-10 Multilink Decomposer Window

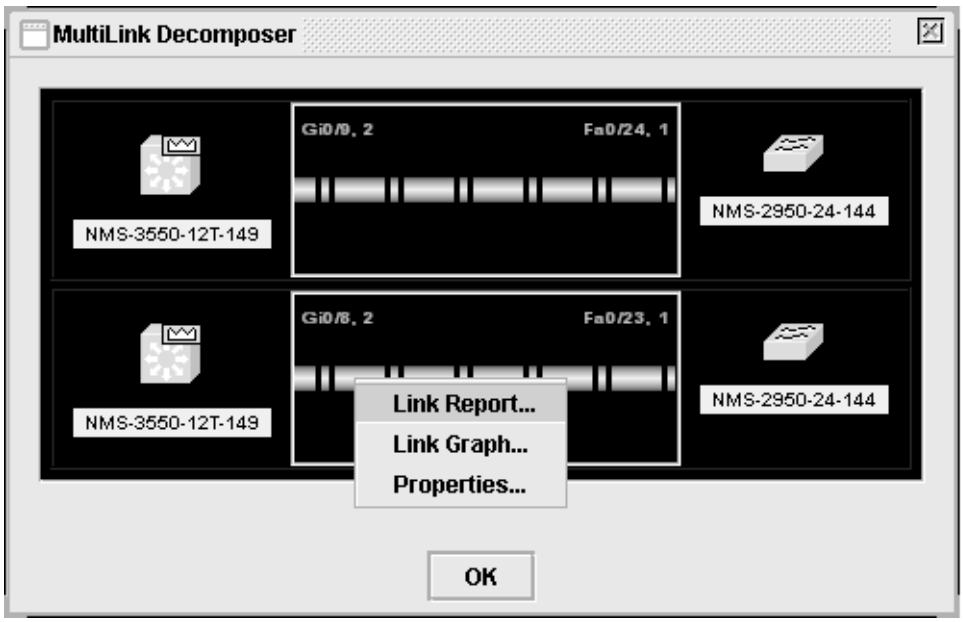

Device Popup Menus

Specific devices in the Topology view display a specific popup menu:

- Cluster (Table 3-15)
- Command switch (Table 3-16)
- Member or standby command switch (Table 3-17)
- Candidate switch with an IP address (Table 3-18)
- Candidate switch without an IP address (Table 3-19)
- Neighboring devices (Table 3-20)

Note The Device Manager option in these popup menus is available in read-only mode on Catalyst 2900 XL and Catalyst 3500 XL switches running Cisco IOS Release 12.0(5)WC2 and later. It is also available on Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 and later and on Catalyst 3550 switch running Cisco IOS Release 12.1(8)EA1 or later. It is not available on the Catalyst 1900 and Catalyst 2820 switches.

To display a device popup menu, click an icon, and right-click.

Table 3-15 Device Popup Menu of a Cluster Icon

Popup Menu Option	Task
Expand cluster	View a cluster-specific topology view.
Properties	Display information about the device and port on either end of the link and the state of the link.

Table 3-16 Device Popup Menu of a Command-Switch Icon

Popup Menu Option	Task
Collapse cluster	View the neighborhood outside a specific cluster.
Host Name[1]	Change the host name of a switch.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use by the switch.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.

Table 3-17 Device Popup Menu of a Member or Standby Command-Switch Icon

Popup Menu Option	Task
Remove from Cluster[1]	Remove a member from the cluster.
Host Name[1]	Change the host name of a switch.
Device Manager[2]	Launch Device Manager for a switch.
Bandwidth Graphs	Display graphs that plot the total bandwidth in use by the switch.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Available only from a cluster-management session.
2. Available from a cluster member switch but not from the command switch.

Table 3-18 Device Popup Menu of a Candidate-Switch Icon (When the Candidate Switch Has an IP Address)

Popup Menu Option	Task
Add to Cluster[1]	Add a candidate to a cluster.
Device Manager[2]	Launch Device Manager for a switch.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.
2. Available from a cluster member switch but not from the command switch.

Table 3-19 Device Popup Menu of a Candidate-Switch Icon (When the Candidate Switch Does Not Have an IP Address)

Popup Menu Option	Task
Add to Cluster[1]	Add a candidate to a cluster.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the "Access Modes in CMS" section on page 3-31.

Table 3-20 Device Popup Menu of a Neighboring-Device Icon

Popup Menu Option	Task
Device Manager[1]	Access the web management interface of the device. **Note** This option is available on Cisco access points, but not on Cisco IP phones, hubs, routers and on *unknown* devices such as some Cisco devices and third-party devices.
Disqualification Code	Display the reason why the device could not join the cluster.
Properties	Display information about the device and port on either end of the link and the state of the link.

1. Available from a cluster member switch but not from the command switch.

Interaction Modes

You can change the interaction mode of CMS to either guide or expert mode. Guide mode steps you through each feature option and provides information about the parameter. Expert mode displays a configuration window in which you configure the feature options.

Guide Mode

Note Guide mode is not available if your switch access level is read-only. For more information about the read-only access mode, see the "Access Modes in CMS" section on page 3-31.

Guide mode is for users who want a step-by-step approach for completing a specific configuration task. This mode is not available for all features. A menu-bar option that has a person icon means that guide mode is available for that option.

When you click **Guide Mode** and then select a menu-bar option that supports guide mode, CMS displays a specific parameter of the feature with information about the parameter field. To configure the feature, you provide the information that CMS requests in each step until you click **Finish** in the last step. Clicking **Cancel** at any time closes and ends the configuration task without applying any changes.

If **Expert Mode** is selected and you want to use guide mode, you must click **Guide Mode** before selecting an option from the menu bar, tool bar, or popup menu. If you change the interaction mode after selecting a configuration option, the mode change does not take effect until you select another configuration option.

Expert Mode

Expert mode is for users who prefer to display all the parameter fields of a feature in a single CMS window. Information about the parameter fields are provided from **Help**.

Wizards

Note Wizards are not available if your switch access level is read-only. For more information about the read-only access mode, see the "Access Modes in CMS" section on page 3-31.

Wizards simplify some configuration tasks on the switch. Similar to the guide mode, wizards provide a step-by-step approach for completing a specific configuration task. Unlike guide mode, a wizard does not prompt you to provide information for all of the feature options. Instead, it prompts you to provide minimal information and then uses the default settings of the remaining options to set up default configurations.

Wizards are not available for all features. A menu-bar option that has *wizard* means that selecting that option launches the wizard for that feature.

Tool Tips

CMS displays a popup message when you move your mouse over these devices:

- A yellow device icon in the cluster tree or in Topology view—A popup displays a fault message, such as that the RPS is faulty or that the switch is unavailable because you are in read-only mode.
- A red device icon in the cluster tree or in Topology view—A popup displays a message that the switch is down.
- A table column heading—A popup displays the full heading.

Online Help

CMS provides comprehensive online help to assist you in understanding and performing configuration and monitoring tasks from the CMS windows (Figure 3-11).

- Feature help, available from the menu bar by selecting **Help** > **Contents**, provides background information and concepts on the features.
- Dialog-specific help, available from **Help** on the CMS windows, provides procedures for performing tasks.
- Index of help topics.
- Glossary of terms used in the online help.

You can send us feedback about the information provided in the online help. Click **Feedback** to display an online form. After completing the form, click **Submit** to send your comments to Cisco. We appreciate and value your comments.

Figure 3-11 Help Contents and Index

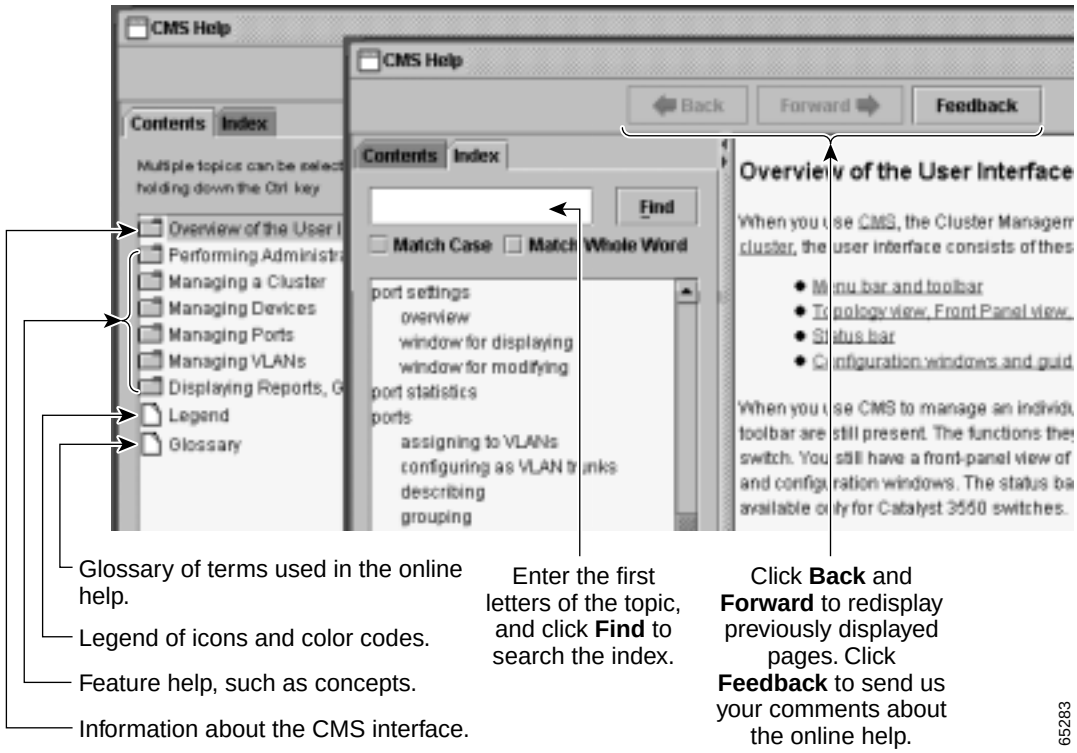

CMS Window Components

CMS windows consistently present configuration information. Figure 3-12 shows the components of a typical CMS window.

Figure 3-12 CMS Window Components

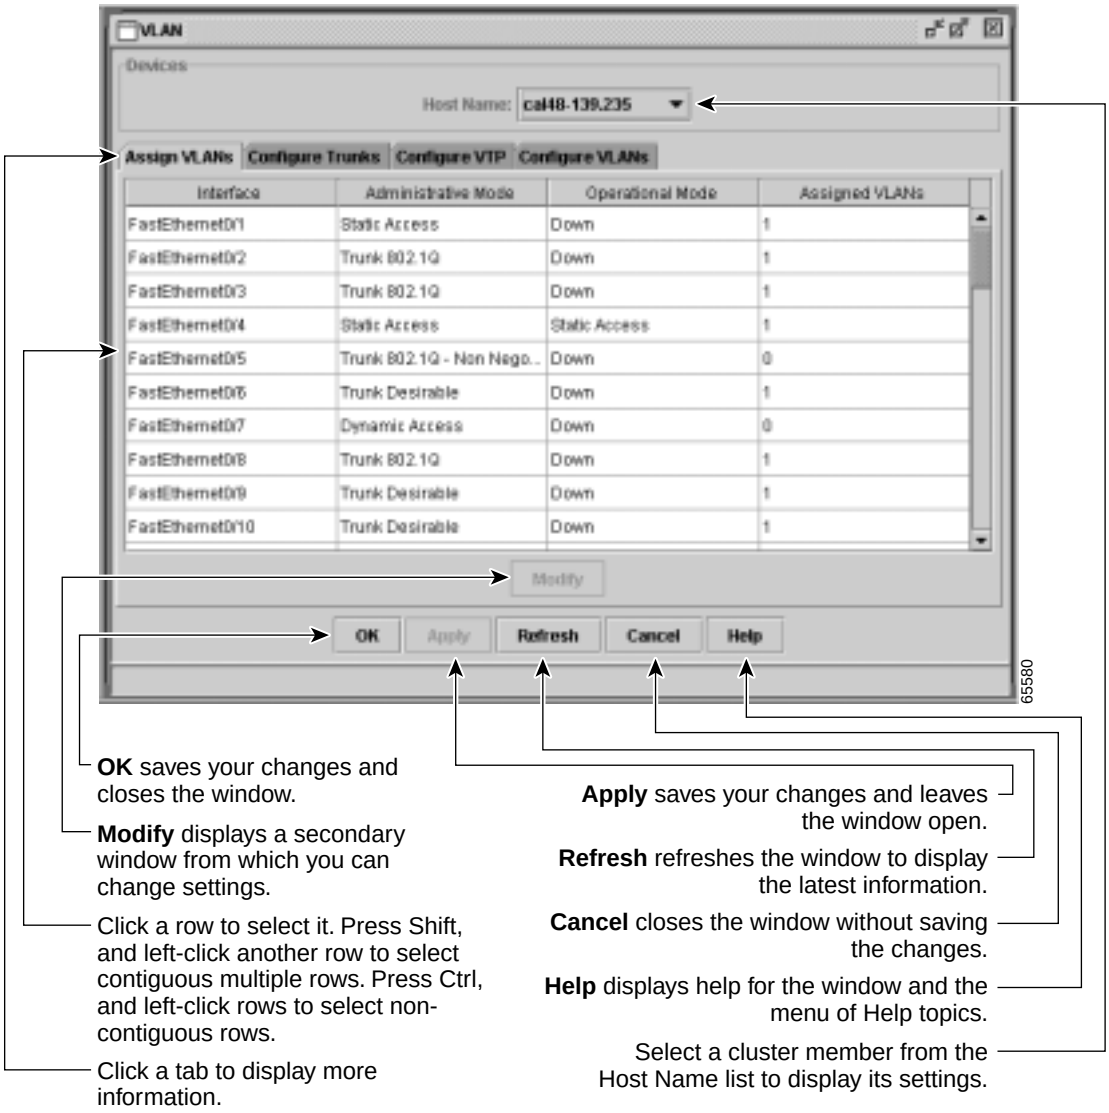
VLAN
Devices
Host Name: ca48-139.235
Assign VLANs
Configure Trunks
Configure VTP
Configure VLANs
Interface
Administrative Mode
Operational Mode
Assigned VLANs
FastEthernet0/1 Static Access Down 1
FastEthernet0/2 Trunk 802.1Q Down 1
FastEthernet0/3 Trunk 802.1Q Down 1
FastEthernet0/4 Static Access Static Access 1
FastEthernet0/5 Trunk 802.1Q - Non Nego... Down 0
FastEthernet0/6 Trunk Desirable Down 1
FastEthernet0/7 Dynamic Access Down 0
FastEthernet0/8 Trunk 802.1Q Down 1
FastEthernet0/9 Trunk Desirable Down 1
FastEthernet0/10 Trunk Desirable Down 1
Modify
OK
Apply
Refresh
Cancel
Help
65580
OK saves your changes and closes the window.
Modify displays a secondary window from which you can change settings.
Click a row to select it. Press Shift, and left-click another row to select contiguous multiple rows. Press Ctrl, and left-click rows to select non-contiguous rows.
Click a tab to display more information.
Apply saves your changes and leaves the window open.
Refresh refreshes the window to display the latest information.
Cancel closes the window without saving the changes.
Help displays help for the window and the menu of Help topics.
Select a cluster member from the Host Name list to display its settings.

Host Name List

To display or change the configuration of a cluster member, you need to select the specific switch from the Host Name drop-down list. The list appears in the configuration window of each feature and lists only the cluster members that support that feature. For example, the Host Name list on the VLAN window does not include Catalyst 1900 and Catalyst 2820 switches even though they are part of the cluster. Similarly, the Host Name list on the LRE Profiles window only lists the LRE switches in the cluster.

Tabs, Lists, and Tables

Some CMS windows have *tabs* that present different sets of information. Tabs are arranged like folder headings across the top of the window. Click the tab to display its information.

Listed information can often be changed by selecting an item from a list. To change the information, select one or more items, and click **Modify**. Changing multiple items is limited to those items that apply to at least one of the selections.

Some CMS windows present information in a table format. You can edit the information in these tables.

Note You can resize the width of the columns to display the column headings, or you can hover your cursor over the heading to display a popup description of the column.

Icons Used in Windows

Some window have icons for sorting information in tables, for showing which cells in a table are editable, and for displaying further information from Cisco.com (Figure 3-13).

Figure 3-13 Window Icons

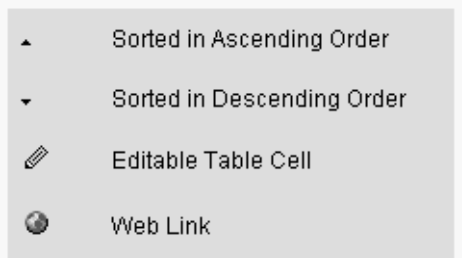

Buttons

These are the most common buttons that you use to change the information in a CMS window:

- OK—Save any changes and close the window. If you made no changes, the window closes. If CMS detects errors in your entry, the window remains open. For more information about error detection, see the "Error Checking" section on page 3-32.
- Apply—Save any changes made in the window and leave the window open. If you made no changes, the Apply button is disabled.
- Refresh—Update the CMS window with the latest status of the device. Unsaved changes are lost.
- Cancel—Do not save any changes made in the window and close the window.
- Help—Display procedures on performing tasks from the window.
- Modify—Display the secondary window for changing information on the selected item or items. You usually select an item from a list or table and click **Modify**.

Accessing CMS

This section assumes the following:

- You know the IP address and password of the command switch or a specific switch. This information is either:
  - Assigned to the switch by following the setup program, as described in the release notes.
  - Changed on the switch by following the information in the "Assigning Switch Information" section on page 4-2 and "Preventing Unauthorized Access to Your Switch" section on page 6-1. Considerations for assigning IP addresses and passwords to a command switch and cluster members are described in the "IP Addresses" section on page 5-15 and "Passwords" section on page 5-16.
- You know your access privilege level to the switch.
- You have referred to the release notes for system requirements and have followed the procedures for installing the required Java plug-ins and configuring your browser.

Caution Copies of the CMS pages you display are saved in your browser memory cache until you exit the browser session. A password is not required to redisplay these pages, including the Cisco Systems Access page. You can access the CLI by clicking **Monitor the router - HTML access to the command line interface** from a cached copy of the Cisco Systems Access page. To prevent unauthorized access to CMS and the CLI, exit your browser to end the browser session.

To access CMS, follow these steps:

Step 1 Enter the switch IP address and your privilege level in the browser **Location** field (Netscape Communicator) or Address field (Microsoft Internet Explorer). For example:

```
http://10.1.126.45:184/level/14/
```

where `10.1.126.45` is the switch IP address, `184` is the HTTP port, and level `14` is the privilege level. You do not need to enter the HTTP port if the switch is using HTTP port 80 (the default) or enter the privilege level if you have read-write access to the switch (privilege level is 15). For information about the HTTP port, see the "HTTP Access to CMS" section on page 3-31. For information about privilege levels, see the "Access Modes in CMS" section on page 3-31.

Step 2 When prompted for a username and password, enter only the switch enable password. CMS prompts you a second time for a username and password. Enter only the enable password again.

If you configure a local username and password, make sure you enable it by using the **ip http authentication** global configuration command. Enter your username and password when prompted.

Step 3 Click **Web Console**.

If you access CMS from a standalone or member switch, Device Manager appears. If you access CMS from a command switch, you can display the Front Panel and Topology views.

Access Modes in CMS

CMS provides two levels of access to the configuration options: read-write access and read-only access. Privilege levels 0 to 15 are supported.

- Privilege level 15 provides you with read-write access to CMS.
- Privilege levels 1 to 14 provide you with read-only access to CMS. Any options in the CMS windows, menu bar, toolbar, and popup menus that change the switch or cluster configuration are not shown in read-only mode.
- Privilege level 0 denies access to CMS.

If you do not include a privilege level when you access CMS, the switch verifies if you have privilege-level 15. If you do not, you are denied access to CMS. If you do have privilege-level 15, you are granted read-write access. Therefore, you do not need to include the privilege level if it is 15. Entering zero denies access to CMS. For more information about privilege levels, see the "Preventing Unauthorized Access to Your Switch" section on page 6-1.

Note

- If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:
  - Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 3550 member switches running Cisco IOS Release 12.1(6)EA1 or earlier

  For more information about this limitation, refer to the Catalyst 3550 release notes.
- These switches do not support read-only mode on CMS:
  - Catalyst 1900 and Catalyst 2820
  - Catalyst 2900 XL switches with 4-MB CPU DRAM

  In read-only mode, these switches appear as unavailable devices and cannot be configured from CMS.

HTTP Access to CMS

CMS uses Hypertext Transfer Protocol (HTTP), which is an in-band form of communication with the switch through any one of its Ethernet ports and that allows switch management from a standard web browser. The default HTTP port is 80.

If you change the HTTP port, you must include the new port number when you enter the IP address in the browser **Location** or **Address** field (for example, http://10.1.126.45:184 where 184 is the new HTTP port number).

Do not disable or otherwise misconfigure the port through which your management station is communicating with the switch. You might want to write down the port number to which you are connected. Changes to the switch IP information should be done with care.

For information about connecting to a switch port, refer to the switch hardware installation guide.

Verifying Your Changes

CMS provides notification cues to help you track and confirm the changes you make.

Change Notification

A green border around a field or table cell means that you made an unsaved change to the field or table cell. Previous information in that field or table cell is displayed in the window status bar. When you save the changes or if you cancel the change, the green border disappears.

Error Checking

A red border around a field means that you entered invalid data in the field. An error message also displays in the window status bar. When you enter valid data in the field, a green border replaces the red border until you either save or cancel the change.

If there is an error in communicating with the switch or if you make an error while performing an action, a popup dialog notifies you about the error.

Saving Your Changes

Note The Save Configuration option is not available if your switch access level is read-only. For more information about the read-only access mode, see the "Access Modes in CMS" section on page 3-31.

Tip As you make cluster configuration changes (except for changes to the Topology view and in the Preferences window), make sure that you periodically save the configuration from the command switch. The configuration is saved on the command and member switches.

The front-panel images and CMS windows always display the *running configuration* of the switch. When you make a configuration change to a switch or switch cluster, the change becomes part of the running configuration. The change *does not* automatically become part of the config.txt file in Flash memory, which is the *startup configuration* used each time the switch restarts. If you do not save your changes to Flash memory, they are lost when the switch restarts.

To save all configuration changes to Flash memory, you must select **Administration > Save Configuration**.

Note Catalyst 1900 and Catalyst 2820 switches automatically save configuration changes to Flash memory as they occur.

Using Different Versions of CMS

When managing switch clusters through CMS, remember that clusters can have a mix of switch models using different IOS releases and that CMS in earlier IOS releases and on different switch platforms might look and function differently from CMS in this IOS release.

When you select **Device > Device Manager** for a cluster member, a new browser session is launched, and the CMS version for that switch is displayed.

Here are examples of how CMS can differ between IOS releases and switch platforms:

- On Catalyst switches running Cisco IOS Release 12.0(5)WC2 or earlier or Cisco IOS Release 12.1(6)EA1 or earlier, the CMS versions in those software releases might appear similar but are not the same as this release. For example, the Topology view in this release is not the same as the Topology view or Cluster View in those earlier software releases.
- CMS on the Catalyst 1900 and Catalyst 2820 switches is referred to as *Switch Manager*. Cluster management options are not available on these switches. This is the earliest version of CMS.

Refer to the documentation specific to the switch and its IOS release for descriptions of the CMS version you are using.

Where to Go Next

The rest of this guide provides information about and CLI procedures for the software features supported in this release. For CMS procedures and window descriptions, refer to the online help.

Refer to the release notes for:

- CMS software requirements
- Procedures for browser configuration
- Procedures for accessing CMS

CHAPTER 4

Assigning the Switch IP Address and Default Gateway

This chapter describes how to create the initial switch configuration (for example, assign the switch IP address and default gateway information) by using a variety of automatic and manual methods. It also describes how to modify the switch startup configuration.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding the Boot Process, page 4-1
- Assigning Switch Information, page 4-2
- Checking and Saving the Running Configuration, page 4-10
- Modifying the Startup Configuration, page 4-12
- Scheduling a Reload of the Software Image, page 4-17

Understanding the Boot Process

Before you can assign switch information (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth), you need to install and power on the switch as described in the hardware installation guide that shipped with your switch.

The normal boot process involves the operation of the boot loader software, which performs these activities:

- Performs low-level CPU initialization. It initializes the CPU registers, which control where physical memory is mapped, its quantity, its speed, and so forth.
- Performs power-on self-test (POST) for the CPU subsystem. It tests the CPU DRAM and the portion of the Flash device that makes up the Flash file system.
- Initializes the Flash file system on the system board.
- Loads a default operating system software image into memory and boots the switch.

The boot loader provides access to the Flash file system before the operating system is loaded. Normally, the boot loader is used only to load, uncompress, and launch the operating system. After the boot loader gives the operating system control of the CPU, the boot loader is not active until the next system reset or power-on.

The boot loader also provides trap-door access into the system if the operating system has problems serious enough that it cannot be used. The trap-door mechanism provides enough access to the system so that if it is necessary, you can format the Flash file system, reinstall the operating system software image by using the XMODEM Protocol, recover from a lost or forgotten password, and finally restart the operating system. For more information, see the "Recovering from Corrupted Software" section on page 27-2 and the "Recovering from a Lost or Forgotten Password" section on page 27-3.

Note On Catalyst 3550 Fast Ethernet switches only, you can disable password recovery. For more information, see the "Disabling Password Recovery" section on page 6-5.

Before you can assign switch information, make sure you have connected a PC or terminal to the console port, and configured the PC or terminal-emulation software baud rate and character format to match those of the switch console port. For more information, refer to the hardware installation guide that shipped with your switch.

Assigning Switch Information

You can assign IP information through the switch setup program, through a Dynamic Host Configuration Protocol (DHCP) server, or manually.

Use the switch setup program if you are a new user and want to be prompted for specific IP information. With this program, you can also configure a host name and an enable secret password. It gives you the option of assigning a Telnet password (to provide security during remote management) and configuring your switch as a command or member switch of a cluster or as a standalone switch. For more information about the setup program, refer to the release notes on Cisco.com.

Use a DHCP server for centralized control and automatic assignment of IP information once the server is configured.

Note If you are using DHCP, do not respond to any of the questions in the setup program until the switch receives the dynamically-assigned IP address and reads the configuration file.

Use the manual method of configuration if you are an experienced user familiar with the switch configuration steps; otherwise, use the setup program described earlier.

This section contains this configuration information:

- Default Switch Information, page 4-3
- Understanding DHCP-Based Autoconfiguration, page 4-3
- Manually Assigning IP Information, page 4-10

Default Switch Information

Table 4-1 shows the default switch information.

Table 4-1 Default Switch Information

Feature	Default Setting
IP address and subnet mask	No IP address or subnet mask are defined.
Default gateway	No default gateway is defined.
Enable secret password	No password is defined.
Host name	The factory-assigned default host name is *Switch*.
Telnet password	No password is defined.
Cluster command switch functionality	Disabled.
Cluster name	No cluster name is defined.

Understanding DHCP-Based Autoconfiguration

The DHCP provides configuration information to Internet hosts and internetworking devices. This protocol consists of two components: one for delivering configuration parameters from a DHCP server to a device and a mechanism for allocating network addresses to devices. DHCP is built on a client-server model, in which designated DHCP servers allocate network addresses and deliver configuration parameters to dynamically configured devices.

During DHCP-based autoconfiguration, your switch (DHCP client) is automatically configured at startup with IP address information and a configuration file.

With DHCP-based autoconfiguration, no DHCP client-side configuration is needed on your switch. However, you need to configure the DHCP server for various lease options associated with IP addresses. If you are using DHCP to relay the configuration file location on the network, you might also need to configure a Trivial File Transfer Protocol (TFTP) server and a Domain Name System (DNS) server.

The DHCP server can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. A relay device forwards broadcast traffic between two directly connected LANs. A router does not forward broadcast packets, but it forwards packets based on the destination IP address in the received packet.

DHCP-based autoconfiguration replaces the BOOTP client functionality on your switch.

DHCP Client Request Process

When you boot your switch, the DHCP client is invoked and automatically requests configuration information from a DHCP server when the configuration file is not present on the switch.

Figure 4-1 shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 4-1 DHCP Client and Server Message Exchange

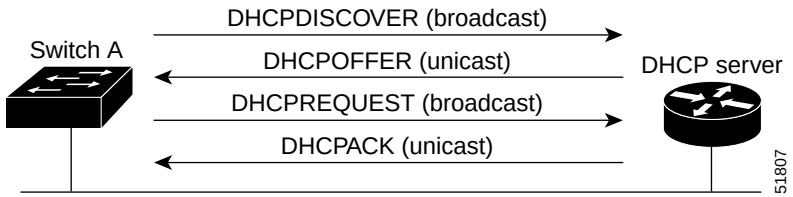

The client, Switch A, broadcasts a DHCPDISCOVER message to locate a DHCP server. The DHCP server offers configuration parameters (such as an IP address, subnet mask, gateway IP address, DNS IP address, a lease for the IP address, and so forth) to the client in a DHCPOFFER unicast message.

In a DHCPREQUEST broadcast message, the client returns a formal request for the offered configuration information to the DHCP server. The formal request is broadcast so that all other DHCP servers that received the DHCPDISCOVER broadcast message from the client can reclaim the IP addresses that they offered to the client.

The DHCP server confirms that the IP address has been allocated to the client by returning a DHCPACK unicast message to the client. With this message, the client and server are bound, and the client uses configuration information received from the server. The amount of information the switch receives depends on how you configure the DHCP server. For more information, see the "Configuring the DHCP Server" section on page 4-5.

If the configuration parameters sent to the client in the DHCPOFFER unicast message are invalid (a configuration error exists), the client returns a DHCPDECLINE broadcast message to the DHCP server.

The DHCP server sends the client a DHCPNAK denial broadcast message, which means that the offered configuration parameters have not been assigned, that an error has occurred during the negotiation of the parameters, or that the client has been slow in responding to the DHCPOFFER message (the DHCP server assigned the parameters to another client).

A DHCP client might receive offers from multiple DHCP or BOOTP servers and can accept any of the offers; however, the client usually accepts the first offer it receives. The offer from the DHCP server is not a guarantee that the IP address will be allocated to the client; however, the server usually reserves the address until the client has had a chance to formally request the address. If the switch accepts replies from a BOOTP server and configures itself, the switch broadcasts, instead of unicasts, TFTP requests to obtain the switch configuration file.

Configuring the DHCP Server

You should configure the DHCP server with reserved leases that are bound to each switch by the switch hardware address.

If you want the switch to receive IP address information, you must configure the DHCP server with these lease options:

- IP address of the client (required)
- Subnet mask of the client (required)
- DNS server IP address (optional)
- Router IP address (default gateway address to be used by the switch) (required)

If you want the switch to receive the configuration file from a TFTP server, you must configure the DHCP server with these lease options:

- TFTP server name (required)
- Boot filename (the name of the configuration file that the client needs) (recommended)
- Host name (optional)

Depending on the settings of the DHCP server, the switch can receive IP address information, the configuration file, or both.

If you do not configure the DHCP server with the lease options described earlier, it replies to client requests with only those parameters that are configured. If the IP address and subnet mask are not in the reply, the switch is not configured. If the router IP address or TFTP server name are not found, the switch might send broadcast, instead of unicast, TFTP requests. Unavailability of other lease options does not affect autoconfiguration.

The DHCP server can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. For more information, see the "Configuring the Relay Device" section on page 4-6.

Configuring the TFTP Server

Based on the DHCP server configuration, the switch attempts to download one or more configuration files from the TFTP server. If you configured the DHCP server to respond to the switch with all the options required for IP connectivity to the TFTP server, and if you configured the DHCP server with a TFTP server name, address, and configuration filename, the switch attempts to download the specified configuration file from the specified TFTP server.

If you did not specify the configuration filename, the TFTP server, or if the configuration file could not be downloaded, the switch attempts to download a configuration file by using various combinations of filenames and TFTP server addresses. The files include the specified configuration filename (if any) and these files: network-config, cisconet.cfg, *hostname*.config, or *hostname*.cfg, where *hostname* is the switch's current hostname. The TFTP server addresses used include the specified TFTP server address (if any) and the broadcast address (255.255.255.255).

For the switch to successfully download a configuration file, the TFTP server must contain one or more configuration files in its base directory. The files can include these files:

- The configuration file named in the DHCP reply (the actual switch configuration file).
- The network-confg or the cisconet.cfg file (known as the default configuration files).
- The router-confg or the ciscortr.cfg file (These files contain commands common to all switches. Normally, if the DHCP and TFTP servers are properly configured, these files are not accessed.)

If you specify the TFTP server name in the DHCP server-lease database, you must also configure the TFTP server name-to-IP-address mapping in the DNS-server database.

If the TFTP server to be used is on a different LAN from the switch, or if it is to be accessed by the switch through the broadcast address (which occurs if the DHCP server response does not contain all the required information described earlier), a relay must be configured to forward the TFTP packets to the TFTP server. For more information, see the "Configuring the Relay Device" section on page 4-6. The preferred solution is to configure the DHCP server with all the required information.

Configuring the DNS

The DHCP server uses the DNS server to resolve the TFTP server name to an IP address. You must configure the TFTP server name-to-IP address map on the DNS server. The TFTP server contains the configuration files for the switch.

You can configure the IP addresses of the DNS servers in the lease database of the DHCP server from where the DHCP replies will retrieve them. You can enter up to two DNS server IP addresses in the lease database.

The DNS server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a router.

Configuring the Relay Device

You must configure a relay device when a switch sends broadcast packets that need to be responded to by a host on a different LAN. Examples of broadcast packets that the switch might send are DHCP, DNS, and in some cases, TFTP packets. You must configure this relay device to forward received broadcast packets on an interface to the destination host.

If the relay device is a Cisco router, enable IP routing (**ip routing** global configuration command), and configure a helper addresses by using the **ip helper-address** interface configuration command.

For example, in Figure 4-2, configure the router interfaces as follows:

On interface 10.0.0.2:

```
router(config-if)# ip helper-address 20.0.0.2
router(config-if)# ip helper-address 20.0.0.3
router(config-if)# ip helper-address 20.0.0.4
```

On interface 20.0.0.1

```
router(config-if)# ip helper-address 10.0.0.1
```

Note If the Catalyst 3550 multilayer switch is acting as the relay device, configure the interface as a routed port. For more information, see the "Routed Ports" section on page 8-4 and the "Configuring Layer 3 Interfaces" section on page 8-22.

Figure 4-2 Relay Device Used in Autoconfiguration

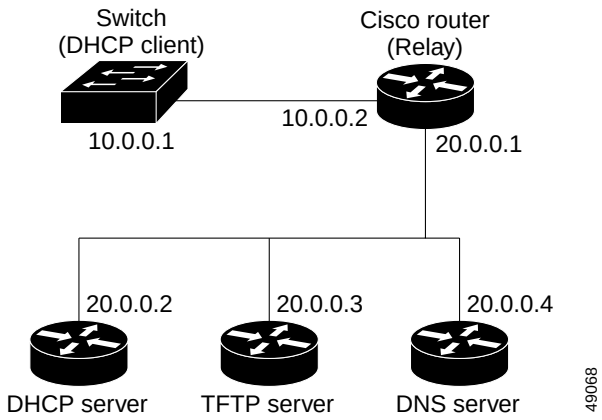

Obtaining Configuration Files

Depending on the availability of the IP address and the configuration filename in the DHCP reserved lease, the switch obtains its configuration information in these ways:

- The IP address and the configuration filename is reserved for the switch and provided in the DHCP reply (one-file read method).

  The switch receives its IP address, subnet mask, TFTP server address, and the configuration filename from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- The IP address and the configuration filename is reserved for the switch, but the TFTP server address is not provided in the DHCP reply (one-file read method).

  The switch receives its IP address, subnet mask, and the configuration filename from the DHCP server. The switch sends a broadcast message to a TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- Only the IP address is reserved for the switch and provided in the DHCP reply. The configuration filename is not provided (two-file read method).

  The switch receives its IP address, subnet mask, and the TFTP server address from the DHCP server. The switch sends a unicast message to the TFTP server to retrieve the network-confg or cisconet.cfg default configuration file. (If the network-confg file cannot be read, the switch reads the cisconet.cfg file.)

  The default configuration file contains the host names-to-IP-address mapping for the switch. The switch fills its host table with the information in the file and obtains its host name. If the host name is not found in the file, the switch uses the host name in the DHCP reply. If the host name is not specified in the DHCP reply, the switch uses the default *Switch* as its host name.

  After obtaining its host name from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its host name (*hostname*-confg or *hostname*.cfg, depending on whether network-confg or cisconet.cfg was read earlier) from the TFTP server. If the cisconet.cfg file is read, the filename of the host is truncated to eight characters.

  If the switch cannot read the network-confg, cisconet.cfg, or the hostname file, it reads the router-confg file. If the switch cannot read the router-confg file, it reads the ciscortr.cfg file.

Note The switch broadcasts TFTP server requests if the TFTP server is not obtained from the DHCP replies, if all attempts to read the configuration file through unicast transmissions fail, or if the TFTP server name cannot be resolved to an IP address.

Example Configuration

Figure 4-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

Figure 4-3 DHCP-Based Autoconfiguration Network Example

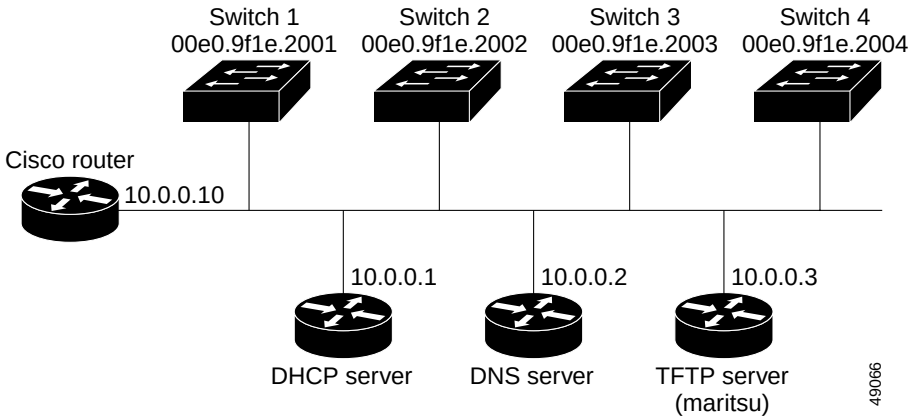

Table 4-2 shows the configuration of the reserved leases on the DHCP server.

Table 4-2 DHCP Server Configuration

	Switch-1	Switch-2	Switch-3	Switch-4
Binding key (hardware address)	00e0.9f1e.2001	00e0.9f1e.2002	00e0.9f1e.2003	00e0.9f1e.2004
IP address	10.0.0.21	10.0.0.22	10.0.0.23	10.0.0.24
Subnet mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Router address	10.0.0.10	10.0.0.10	10.0.0.10	10.0.0.10
DNS server address	10.0.0.2	10.0.0.2	10.0.0.2	10.0.0.2
TFTP server name	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3
Boot filename (configuration file) (optional)	switch1-confg	switch2-confg	switch3-confg	switch4-confg
Host name (optional)	switch1	switch2	switch3	switch4

DNS Server Configuration

The DNS server maps the TFTP server name *maritsu* to IP address 10.0.0.3.

TFTP Server Configuration (on UNIX)

The TFTP server base directory is set to /tftpserver/work/. This directory contains the network-confg file used in the two-file read method. This file contains the host name to be assigned to the switch based on its IP address. The base directory also contains a configuration file for each switch (*switch1-confg*, *switch2-confg*, and so forth) as shown in this display:

```
prompt> cd /tftpserver/work/
prompt> ls
network-confg
switch1-confg
switch2-confg
switch3-confg
switch4-confg
prompt> cat network-confg
ip host switch1 10.0.0.21
ip host switch2 10.0.0.22
ip host switch3 10.0.0.23
ip host switch4 10.0.0.24
```

DHCP Client Configuration

No configuration file is present on Switch 1 through Switch 4.

Configuration Explanation

In Figure 4-3, Switch 1 reads its configuration file as follows:

- It obtains its IP address 10.0.0.21 from the DHCP server.
- If no configuration filename is given in the DHCP server reply, Switch 1 reads the network-confg file from the base directory of the TFTP server.
- It adds the contents of the network-confg file to its host table.
- It reads its host table by indexing its IP address 10.0.0.21 to its host name (switch1).
- It reads the configuration file that corresponds to its host name; for example, it reads *switch1-confg* from the TFTP server.

Switches 2 through 4 retrieve their configuration files and IP addresses in the same way.

Manually Assigning IP Information

Beginning in privileged EXEC mode, follow these steps to manually assign IP information to multiple switched virtual interfaces (SVIs) or ports:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface vlan** *vlan-id*	Enter interface configuration mode, and enter the VLAN to which the IP information is assigned. The range is 1 to 1001.
Step 3	**ip address** *ip-address subnet-mask*	Enter the IP address and subnet mask.
Step 4	**exit**	Return to global configuration mode.
Step 5	**ip default-gateway** *ip-address*	Enter the IP address of the next-hop router interface that is directly connected to the switch where a default gateway is being configured. The default gateway receives IP packets with unresolved destination IP addresses from the switch. Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate. **Note** When your switch is configured to route with IP, it does not need to have a default gateway set.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the switch IP address, use the **no ip address** interface configuration command. If you are removing the address through a Telnet session, your connection to the switch will be lost. To remove the default gateway address, use the **no ip default-gateway** global configuration command.

For information on setting the switch system name, protecting access to privileged EXEC commands, and setting time and calendar services, see Chapter 6, “Administering the Switch.”

Checking and Saving the Running Configuration

You can check the configuration settings you entered or changes you made by entering this privileged EXEC command:

```
Switch# show running-config

Building configuration...

Current configuration: 1363 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
enable secret 5 $1$ej9.$DMUvAUnZOAmvmgqBEzIxE0
```

```
!
interface GigabitEthernet0/1
 no switchport
 ip address 172.20.137.50 255.255.255.0
!
interface GigabitEthernet0/2
!
interface GigabitEthernet0/3
!
interface GigabitEthernet0/4
!
interface GigabitEthernet0/5
!
interface GigabitEthernet0/6
!
interface GigabitEthernet0/7
!
interface GigabitEthernet0/8
!
interface GigabitEthernet0/9
!
interface GigabitEthernet0/10
!
interface GigabitEthernet0/11
!
interface GigabitEthernet0/12
...!
interface VLAN1
 ip address 172.20.137.50 255.255.255.0
 no ip directed-broadcast
 ip nat outside
!
ip default-gateway 172.20.137.1 !
!
snmp-server community private RW
snmp-server community public RO
snmp-server community private@es0 RW
snmp-server community public@es0 RO
snmp-server chassis-id 0x12
!
end
```

To store the configuration or changes you have made to your startup configuration in Flash memory, enter this privileged EXEC command:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
```

This command saves the configuration settings that you made. If you fail to do this, your configuration will be lost the next time you reload the system. To display information stored in the NVRAM section of Flash memory, use the **show startup-config** or **more startup-config** privileged EXEC command.

For more information about alternative locations to copy the configuration file, see Appendix B, “Working with the IOS File System, Configuration Files, and Software Images.”

Modifying the Startup Configuration

This section describes how to modify the switch startup configuration. It contains this configuration information:

- Default Boot Configuration, page 4-12
- Automatically Downloading a Configuration File, page 4-12
- Booting Manually, page 4-13
- Booting a Specific Software Image, page 4-14
- Controlling Environment Variables, page 4-15

Default Boot Configuration

Table 4-3 shows the default boot configuration.

Table 4-3 Default Boot Configuration

Feature	Default Setting
Operating system software image	The switch attempts to automatically boot the system using information in the BOOT environment variable. If the variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the Flash file system. The IOS image is stored in a directory that has the same name as the image file (excluding the .bin extension). In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory.
Configuration file	Configured switches use the *config.text* file stored on the system board in Flash memory. A new switch has no configuration file.

Automatically Downloading a Configuration File

You can automatically download a configuration file to your switch by using the DHCP-based autoconfiguration feature. For more information, see the "Understanding DHCP-Based Autoconfiguration" section on page 4-3.

Specifying the Filename to Read and Write the System Configuration

By default, the IOS software uses the file *config.text* to read and write a nonvolatile copy of the system configuration. However, you can specify a different filename, which will be loaded during the next boot cycle.

Beginning in privileged EXEC mode, follow these steps to specify a different configuration filename:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**boot config-file flash:**/*file-url*	Specify the configuration file to load during the next boot cycle. For *file-url*, specify the path (directory) and the configuration filename. Filenames and directory names are case sensitive.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show boot**	Verify your entries. The **boot config-file** global configuration command changes the setting of the CONFIG_FILE environment variable.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot config-file** global configuration command.

Booting Manually

By default, the switch automatically boots; however, you can configure it to manually boot.

Beginning in privileged EXEC mode, follow these steps to configure the switch to manually boot during the next boot cycle:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**boot manual**	Enable the switch to manually boot during the next boot cycle.
Step 3	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 4	**show boot**	Verify your entries. The **boot manual** global command changes the setting of the MANUAL_BOOT environment variable. The next time you reboot the system, the switch is in boot loader mode, shown by the *switch:* prompt. To boot the system, use the **boot** *filesystem***:**/*file-url* boot loader command. • For *filesystem***:**, use **flash:** for the system board Flash device. • For *file-url*, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable manual booting, use the **no boot manual** global configuration command.

Booting a Specific Software Image

By default, the switch attempts to automatically boot the system using information in the BOOT environment variable. If this variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the Flash file system. In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory. However, you can specify a specific image to boot.

Beginning in privileged EXEC mode, follow these steps to configure the switch to boot a specific image during the next boot cycle:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**boot system** *filesystem***:**/*file-url*	Configure the switch to boot a specific image in Flash memory during the next boot cycle. • For *filesystem***:**, use **flash:** for the system board Flash device. • For *file-url*, specify the path (directory) and the name of the bootable image. Filenames and directory names are case sensitive.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show boot**	Verify your entries. The **boot system** global command changes the setting of the BOOT environment variable. During the next boot cycle, the switch attempts to automatically boot the system using information in the BOOT environment variable.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot system** global configuration command.

Controlling Environment Variables

With a normally operating switch, you enter the boot loader mode only through a switch console connection configured for 9600 bps. Unplug the switch power cord and press the switch **Mode** button while reconnecting the power cord. You can release the **Mode** button a second or two after the LED above port 1X goes off. Then the boot loader *switch:* prompt is displayed.

The switch boot loader software provides support for nonvolatile environment variables, which can be used to control how the boot loader, or any other software running on the system, behaves. Boot loader environment variables are similar to environment variables that can be set on UNIX or DOS systems.

Environment variables that have values are stored in the Flash file system in various files as shown in Table 4-4.

Table 4-4 Environment Variables Storage Location

Environment Variable	Location (file system:filename)
BAUD, ENABLE_BREAK, CONFIG_BUFSIZE, CONFIG_FILE, MANUAL_BOOT, PS1	flash:env_vars
BOOT, BOOTHLPR, HELPER, HELPER_CONFIG_FILE	flash:system_env_vars

Each line in these files contains an environment variable name and an equal sign followed by the value of the variable. A variable has no value if it is not listed in this file; it has a value if it is listed in the file even if the value is a null string. A variable that is set to a null string (for example, " ") is a variable with a value. Many environment variables are predefined and have default values.

Environment variables store two kinds of data:

- Data that controls code, which does not read the IOS configuration file. For example, the name of a boot loader helper file, which extends or patches the functionality of the boot loader can be stored as an environment variable.
- Data that controls code, which is responsible for reading the IOS configuration file. For example, the name of the IOS configuration file can be stored as an environment variable.

You can change the settings of the environment variables by accessing the boot loader or by using IOS commands. Under normal circumstances, it is not necessary to alter the setting of the environment variables.

Note For complete syntax and usage information for the boot loader commands and environment variables, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Table 4-5 describes the function of the most common environment variables.

Table 4-5 Environment Variables

Variable	Boot Loader Command	IOS Global Configuration Command
MANUAL_BOOT	**set MANUAL_BOOT yes** Determines whether the switch automatically or manually boots. Valid values are 1, yes, 0, and no. If it is set to no or 0, the boot loader attempts to automatically boot the system. If it is set to anything else, you must manually boot the switch from the boot loader mode.	**boot manual** Enables manually booting the switch during the next boot cycle and changes the setting of the MANUAL_BOOT environment variable. The next time you reboot the system, the switch is in boot loader mode. To boot the system, use the boot loader **boot flash:***filesystem***:/***file-url* command, and specify the name of the bootable image.
BOOT	**set BOOT** *filesystem***:/***file-url* ... A semicolon-separated list of executable files to try to load and execute when automatically booting. If the BOOT environment variable is not set, the system attempts to load and execute the first executable image it can find by using a recursive, depth-first search through the Flash file system. If the BOOT variable is set but the specified images cannot be loaded, the system attempts to boot the first bootable file that it can find in the Flash file system.	**boot system** *filesystem***:/***file-url* Specifies the IOS image to load during the next boot cycle. This command changes the setting of the BOOT environment variable.
CONFIG_FILE	**set CONFIG_FILE flash:/***file-url* Changes the filename that IOS uses to read and write a nonvolatile copy of the system configuration.	**boot config-file flash:/***file-url* Specifies the filename that IOS uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.
CONFIG_BUFSIZE	**set CONFIG_BUFSIZE** *size* Changes the buffer size that IOS uses to hold a copy of the configuration file in memory. The configuration file cannot be larger than the buffer size allocation. The range is from 4096 to 524288 bytes.	**boot buffersize** *size* Specifies the size of the file system-simulated NVRAM in Flash memory. The buffer holds a copy of the configuration file in memory. This command changes the setting of the CONFIG_BUFSIZE environment variable. You must reload the switch by using the **reload** privileged EXEC command for this command to take effect.

Scheduling a Reload of the Software Image

You can schedule a reload of the software image to occur on the switch at a later time (for example, late at night or during the weekend when the switch is used less), or you can synchronize a reload network-wide (for example, to perform a software upgrade on all switches in the network).

Note A scheduled reload must take place within approximately 24 days.

Configuring a Scheduled Reload

To configure your switch to reload the software image at a later time, use one of these commands in privileged EXEC mode:

- **reload in** [*hh***:**]*mm* [*text*]

  This command schedules a reload of the software to take affect in the specified minutes or hours and minutes. The reload must take place within approximately 24 days. You can specify the reason for the reload in a string up to 255 characters in length.

- **reload at** *hh***:***mm* [*month day* | *day month*] [*text*]

  This command schedules a reload of the software to take place at the specified time (using a 24-hour clock). If you specify the month and day, the reload is scheduled to take place at the specified time and date. If you do not specify the month and day, the reload takes place at the specified time on the current day (if the specified time is later than the current time) or on the next day (if the specified time is earlier than the current time). Specifying 00:00 schedules the reload for midnight.

  Note Use the **at** keyword only if the switch system clock has been set (through Network Time Protocol (NTP), the hardware calendar, or manually). The time is relative to the configured time zone on the switch. To schedule reloads across several switches to occur simultaneously, the time on each switch must be synchronized with NTP.

The **reload** command halts the system. If the system is not set to manually boot, it reboots itself. Use the **reload** command after you save the switch configuration information to the startup configuration (**copy running-config startup-config**).

If your switch is configured for manual booting, do not reload it from a virtual terminal. This restriction prevents the switch from entering the boot loader mode and thereby taking it from the remote user's control.

If you modify your configuration file, the switch prompts you to save the configuration before reloading. During the save operation, the system requests whether you want to proceed with the save if the CONFIG_FILE environment variable points to a startup configuration file that no longer exists. If you proceed in this situation, the system enters setup mode upon reload.

This example shows how to reload the software on the switch on the current day at 7:30 p.m:

```
Switch# reload at 19:30
Reload scheduled for 19:30:00 UTC Wed Jun 5 1996 (in 2 hours and 25 minutes)
Proceed with reload? [confirm]
```

This example shows how to reload the software on the switch at a future time:

```
Switch# reload at 02:00 jun 20
Reload scheduled for 02:00:00 UTC Thu Jun 20 1996 (in 344 hours and 53 minutes)
Proceed with reload? [confirm]
```

To cancel a previously scheduled reload, use the **reload cancel** privileged EXEC command.

Displaying Scheduled Reload Information

To display information about a previously scheduled reload or to determine if a reload has been scheduled on the switch, use the **show reload** privileged EXEC command.

It displays reload information including the time the reload is scheduled to occur and the reason for the reload (if it was specified when the reload was scheduled).

Clustering Switches

This chapter provides these topics to help you get started with switch clustering:

- Understanding Switch Clusters, page 5-2
- Planning a Switch Cluster, page 5-4
- Creating a Switch Cluster, page 5-18
- Using the CLI to Manage Switch Clusters, page 5-25
- Using SNMP to Manage Switch Clusters, page 5-26

Configuring switch clusters is more easily done from the Cluster Management Suite (CMS) web-based interface than through the command-line interface (CLI). Therefore, information in this chapter focuses on using CMS to create a cluster. See Chapter 3, "Getting Started with CMS," for additional information about switch clusters and the clustering options. For complete procedures on using CMS to configure switch clusters, refer to the online help.

For the CLI cluster commands, refer to the switch command reference.

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, and for the required software versions and browser and Java plug-in configurations.

Note This chapter focuses on Catalyst 3550 switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable Catalyst switches, but it does not provide complete descriptions of the cluster features for these other switches. For complete cluster information for a specific Catalyst platform, refer to the software configuration guide for that switch.

Understanding Switch Clusters

A switch cluster is a group of connected Catalyst switches that are managed as a single entity. In a switch cluster, 1 switch must be the *command switch* and up to 15 switches can be *member switches*. The total number of switches in a cluster cannot exceed 16 switches. The command switch is the single point of access used to configure, manage, and monitor the member switches. Cluster members can belong to only one cluster at a time.

The benefits of clustering switches include:

- Management of Catalyst switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 (if your cluster is using a Catalyst 3550 multilayer switch as a Layer 3 router between the Layer 2 switches in the cluster) network. Cluster members are connected through at least one VLAN in common with the command switch according to the connectivity guidelines described in the "Automatic Discovery of Cluster Candidates and Members" section on page 5-4. This section includes some management VLAN considerations for the Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.
- Command-switch redundancy if a command switch fails. One or more switches can be designated as *standby command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby command switches.
- Management of a variety of Catalyst switches through a single IP address. This conserves on IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the command switch IP address.

For other clustering benefits, see the "Advantages of Using CMS and Clustering Switches" section on page 1-6.

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, and the required software versions.

Command Switch Characteristics

A Catalyst 3550 command switch must meet these requirements:

- It is running Cisco IOS Release 12.1(4)EA1 or later.
- It has an IP address.
- It has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- It is not a command or member switch of another cluster.

Note

- We strongly recommend that the highest-end, command-capable switch in the cluster be the command switch:
- If your switch cluster has a Catalyst 3550 switch, that switch should be the command switch.
- If your switch cluster has Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches, the Catalyst 2950 should be the command switch.
- If your switch cluster has Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, and Catalyst 3500 XL switches, either the Catalyst 2900 XL or Catalyst 3500 XL should be the command switch.

Standby Command Switch Characteristics

A Catalyst 3550 standby command switch must meet these requirements:

- It is running 12.1(4)EA1 or later.
- It has an IP address.
- It has CDP version 2 enabled.
- It is connected to the command switch and all other standby command switches through at least one common VLAN.
- It is redundantly connected to the cluster so that connectivity to member switches is maintained.
- It is not a command or member switch of another cluster.

Note

- When the command switch is a Catalyst 3550 switch, all standby command switches must be Catalyst 3550 switches or Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the command switch is a Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the Catalyst 2950 command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches also running Cisco IOS Release 12.0(5)WC2 or earlier: Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches.

Candidate and Member Switches Characteristics

Candidate switches are cluster-capable switches that have not yet been added to a cluster. Member switches are switches that have actually been added to a switch cluster. Although not required, a candidate or member switch can have its own IP address and password (for related considerations, see the "IP Addresses" section on page 5-15 and "Passwords" section on page 5-16).

To join a cluster, a candidate switch must meet these requirements:

- It is running cluster-capable software.
- It has CDP version 2 enabled.
- It is not a command or member switch of another cluster.
- It is connected to the command switch through at least one common VLAN.
- (If a cluster standby group exists) It is connected to every standby command switch through at least one common VLAN. The VLAN to each standby command switch can be different.

Note Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL member switches must be connected through their management VLAN to the command switch and standby command switches. This requirement does not apply to Catalyst 3550 switches, which can connect to any VLAN in common with the command switch.

Planning a Switch Cluster

Anticipating conflicts and compatibility issues is a high priority when you manage several switches through a cluster. This section describes these guidelines, requirements, and caveats that you should understand before you create the cluster:

- Automatic Discovery of Cluster Candidates and Members, page 5-4
- HSRP and Standby Command Switches, page 5-12
- IP Addresses, page 5-15
- Host Names, page 5-16
- Passwords, page 5-16
- SNMP Community Strings, page 5-16
- TACACS+, page 5-17
- Access Modes in CMS, page 5-17
- LRE Profiles, page 5-17
- Availability of Switch-Specific Features in Switch Clusters, page 5-18

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, and for the required software versions and browser and Java plug-in configurations.

Automatic Discovery of Cluster Candidates and Members

The command switch uses Cisco Discovery Protocol (CDP) to discover member switches, candidate switches, neighboring switch clusters, and edge devices across multiple VLANs and in star or cascaded topologies.

Note Do not disable CDP on the command switch, on cluster members, or on any cluster-capable switches that you might want a command switch to discover. For more information about CDP, see Chapter 13, "Configuring CDP."

Following these connectivity guidelines ensures automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices:

- Discovery through CDP Hops, page 5-5
- Discovery through Non-CDP-Capable and Noncluster-Capable Devices, page 5-6
- Discovery through Different VLANs, page 5-7
- Discovery through Different Management VLANs, page 5-9
- Discovery through the Same Management VLAN, page 5-8
- Discovery through Routed Ports, page 5-10
- Discovery of Newly Installed Switches, page 5-11

Discovery through CDP Hops

By using CDP, a command switch can discover switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last member switches are connected to the cluster (for example, the command switch and member switches 8, 9, and 10 in Figure 5-1 are at the edge of the cluster).

You can set the number of hops the command switch searches for candidate and member switches by selecting **Cluster** > **Hop Count**. When new candidate switches are added to the network, the command switch discovers them and adds them to the list of candidate switches.

Figure 5-1 shows a switch cluster with candidate switches. The command switch has ports assigned to VLANs 16 and 62. The CDP hop count is three. The command switch discovers switches 11, 12, 13, and 14 because they are within 3 hops from the edge of the cluster. It does not discover switch 15 because it is 4 hops from the edge of the cluster.

Figure 5-1 Discovery through CDP Hops

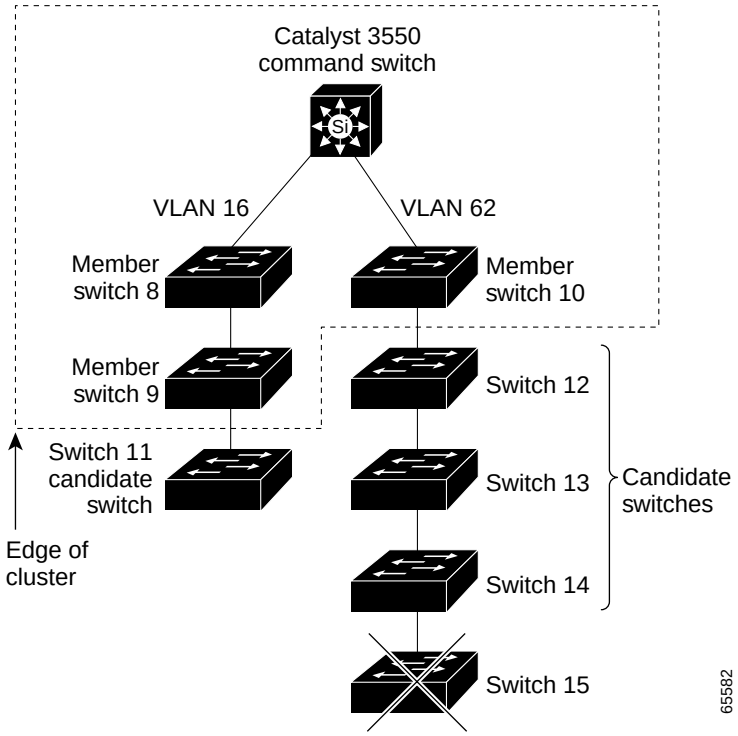

Discovery through Non-CDP-Capable and Noncluster-Capable Devices

If a command switch is connected to a *non-CDP-capable third-party hub* (such as a non-Cisco hub), it can discover cluster-enabled devices connected to that third-party hub. However, if the command switch is connected to a *noncluster-capable Cisco device*, it cannot discover a cluster-enabled device connected beyond the noncluster-capable Cisco device.

Figure 5-2 shows that the command switch discovers the Catalyst 3500 XL switch, which is connected to a third-party hub. However, the command switch does not discover the Catalyst 3550 switch that is connected to a Catalyst 5000 switch.

Figure 5-2 Discovery through Non-CDP-Capable and Noncluster-Capable Devices

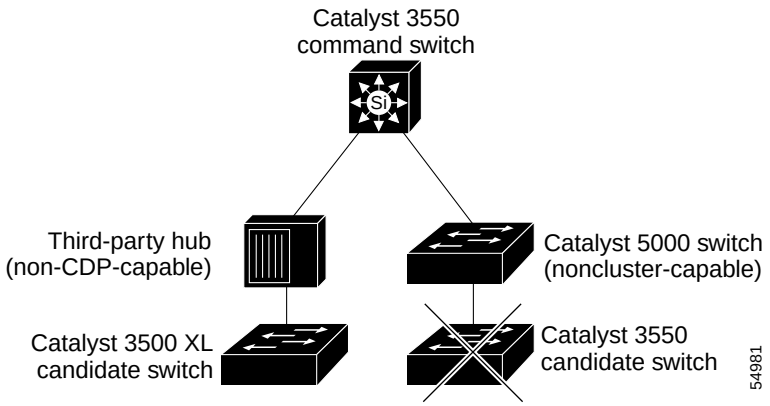

Discovery through Different VLANs

A cluster can have Catalyst 3550 member switches configured with different VLANs. However, each member switch must be connected through at least one VLAN in common with the command switch. The command switch in Figure 5-3 has ports assigned to VLANs 9, 16, and 62 and therefore discovers the switches in those VLANs. It does not discover the switch in VLAN 50. It also does not discover the switch in VLAN 16 in the first column because the command switch has no VLAN connectivity to it.

For more information about VLANs, see Chapter 9, "Creating and Maintaining VLANs." For information about discovery through management VLANs, see the "Discovery through the Same Management VLAN" section on page 5-8 and "Discovery through Different Management VLANs" section on page 5-9.

Figure 5-3 Discovery through Different VLANs

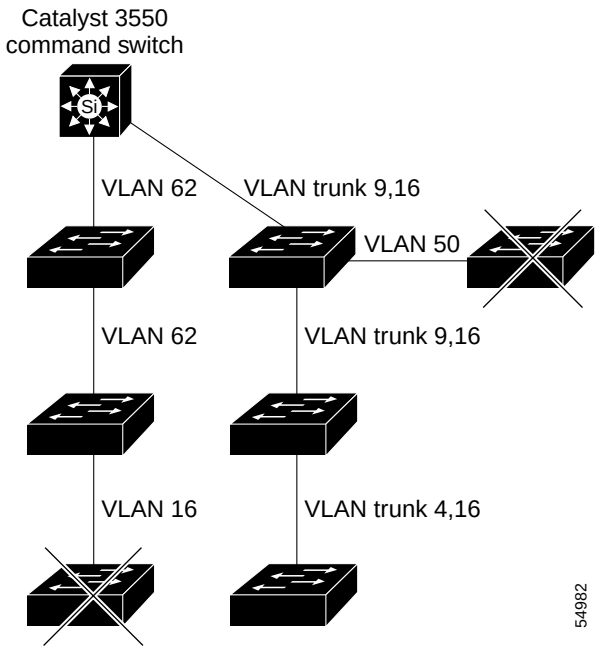

Discovery through the Same Management VLAN

When the cluster has a Catalyst 2900 XL, Catalyst 2950, or Catalyst 3500 XL command switch, all cluster members must connect to it through the command-switch management VLAN, which is VLAN 1 by default. If the cluster members include Catalyst 3550 switches, these member switches must also be connected to the command-switch management VLAN. For information about management VLANs on these switches, refer to the software configuration guide for that specific switch.

To avoid this limitation, assign a Catalyst 3550 switch as the command switch whenever possible. A Catalyst 3550 command switch can manage these cluster members even if they belong to different management VLANs (see the "Discovery through Different Management VLANs" section on page 5-9).

The command switch in Figure 5-4 has ports assigned to management VLAN 9. It discovers all but these switches:

- Switches 7 and 10 because their management VLAN (VLAN 4) is different from the command-switch management VLAN (VLAN 9)
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 5-4 Discovery through the Same Management VLAN

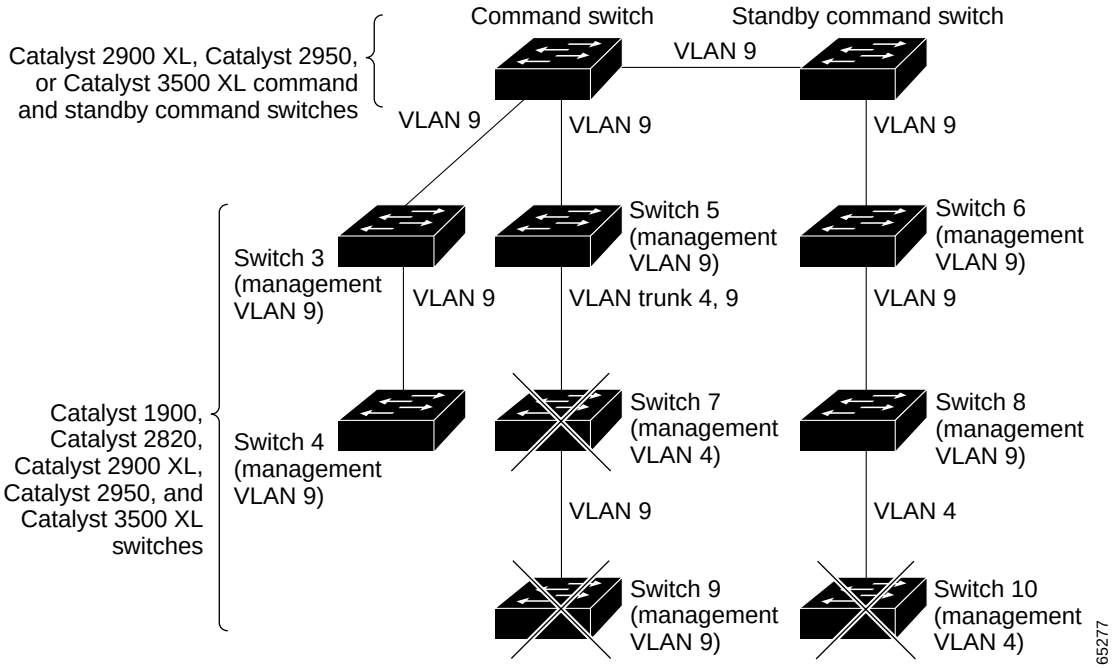
Command switch
Standby command switch
Catalyst 2900 XL, Catalyst 2950, or Catalyst 3500 XL command and standby command switches
VLAN 9
VLAN 9
VLAN 9
VLAN 9
Switch 5 (management VLAN 9)
Switch 6 (management VLAN 9)
Switch 3 (management VLAN 9)
VLAN 9
VLAN trunk 4, 9
VLAN 9
Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches
Switch 4 (management VLAN 9)
Switch 7 (management VLAN 4)
Switch 8 (management VLAN 9)
VLAN 9
VLAN 4
Switch 9 (management VLAN 9)
Switch 10 (management VLAN 4)
65277

Discovery through Different Management VLANs

We strongly recommend that a Catalyst 3550 switch be the command switch when the cluster has Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL member switches. These member switches must connect to each other and to a Catalyst 3550 command switch through their management VLAN, which is VLAN 1 by default. When these member switches are directly connected to a Catalyst 3550 command switch, the management VLAN of these member switches can be different.

For information about discovery through the same management VLAN on these switches, see the "Discovery through the Same Management VLAN" section on page 5-8. For information about management VLANs on these switches, refer to the software configuration guide for that specific switch.

The command switch in Figure 5-5 has ports assigned to VLANs 9, 16, and 62. It discovers all the switches in the different management VLANs except these:

- Switches 7 and 10 because their management VLAN (VLAN 4) is different from the command-switch management VLAN (VLAN 9)
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 5-5 Discovery through Different Management VLANs

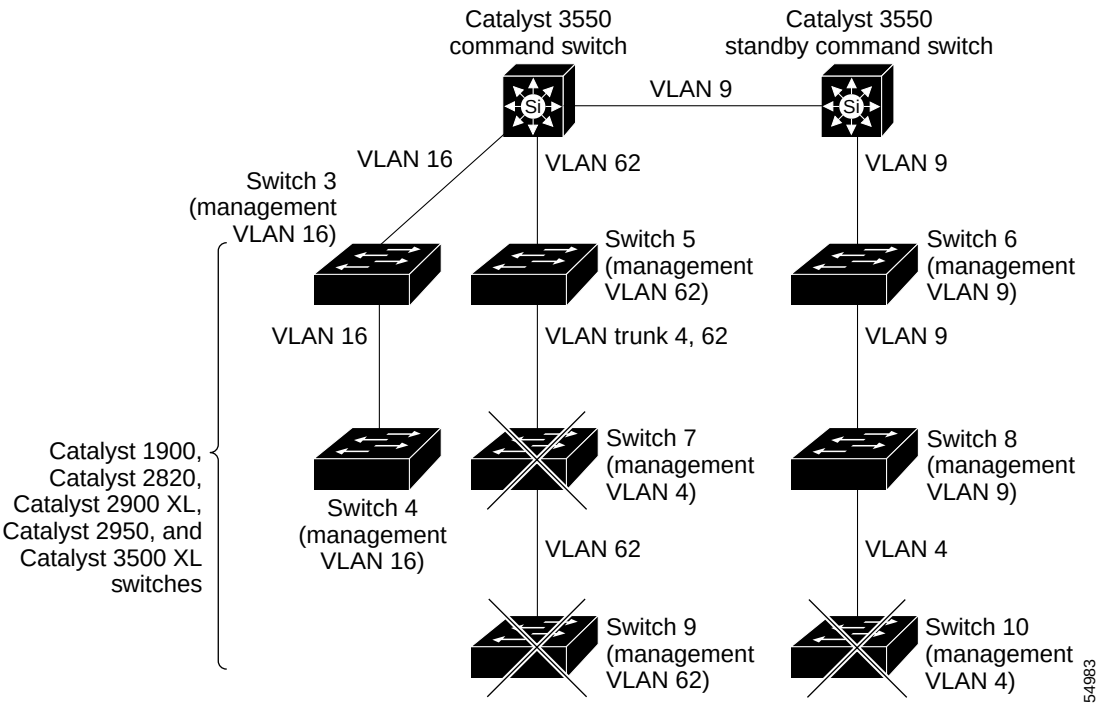

Discovery through Routed Ports

If the command switch has a routed port (RP) configured, it discovers only candidate and member switches in the *same* VLAN as the routed port. For more information about routed ports, see the "Routed Ports" section on page 8-4.

The command switch in Figure 5-6 can discover the switches in VLANs 9 and 62 but not the switch in VLAN 4. If the routed port path between the command switch and member switch 7 is lost, connectivity with member switch 7 is maintained because of the redundant path through VLAN 9.

Figure 5-6 Discovery through Routed Ports

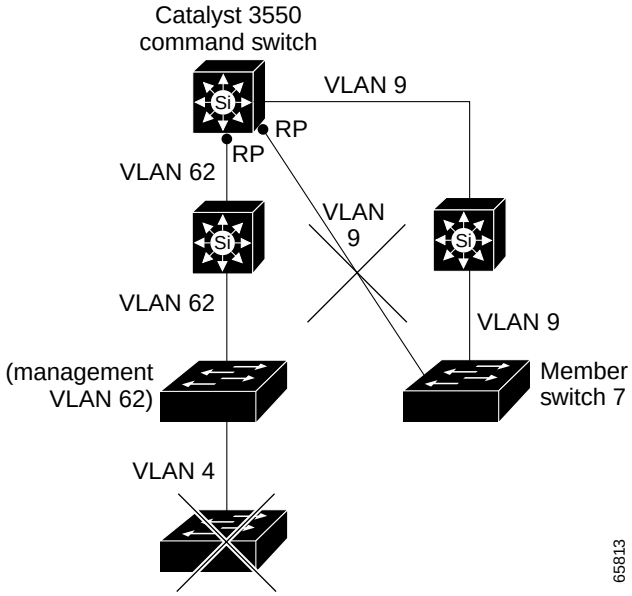

Discovery of Newly Installed Switches

A new, out-of-the-box switch is set with the default VLAN, VLAN 1. By default, all access ports on the new switch are assigned to VLAN 1.

To add a new switch to a cluster, it must be connected to the cluster through an access port. When the new switch joins a cluster, its default VLAN changes to the VLAN of the immediately upstream neighbor. The new switch also configures its access port to belong to the VLAN of the immediately upstream neighbor.

An access port (AP) carries the traffic of and belongs to only one VLAN. For more information about access ports, see the "Access Ports" section on page 8-2.

For example, the command switch in Figure 5-7 belongs to VLAN 9 and VLAN 16.

- A new Catalyst 3550 switch automatically configures the access port to belong to the immediately upstream VLAN, VLAN 9.
- A new Catalyst 2950 switch configures the access port to belong to the upstream VLAN, VLAN 16. The management VLAN of the Catalyst 2950 switch becomes VLAN 16.

Figure 5-7 Discovery of Newly Installed Switches

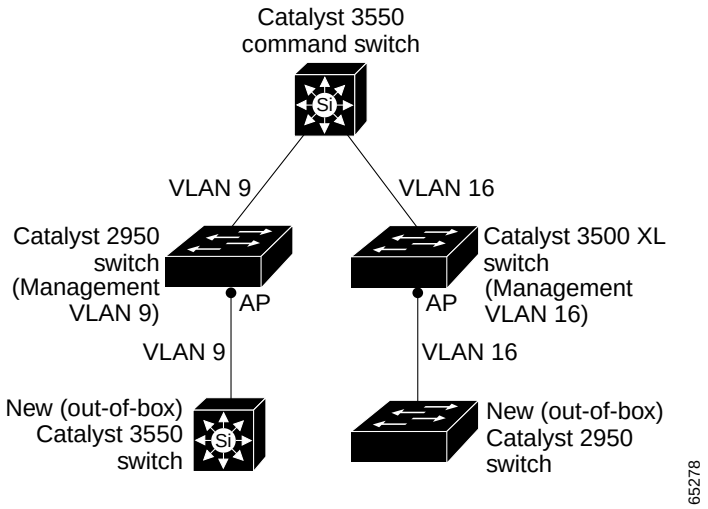

HSRP and Standby Command Switches

The switch supports Hot Standby Router Protocol (HSRP) so that you can configure a group of standby command switches. Because a command switch manages the forwarding of all communication and configuration information to all the member switches, we strongly recommend that you configure a cluster standby command switch to take over if the primary command switch fails.

A *cluster standby group* is a group of command-capable switches that meet the requirements described in the "Standby Command Switch Characteristics" section on page 5-3. Only one cluster standby group can be assigned per cluster.

Note

- When the command switch is a Catalyst 3550 switch, all standby command switches must be Catalyst 3550 switches or Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the command switch is a Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the Catalyst 2950 command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches also running Cisco IOS Release 12.0(5)WC2 or earlier: Catalyst 2900 XL, Catalyst 2950, andCatalyst 3500 XL switches.

Note The cluster standby group is an HSRP group. Disabling HSRP disables the cluster standby group.

The switches in the cluster standby group are ranked according to HSRP priorities. The switch with the highest priority in the group is the *active command switch* (AC). The switch with the next highest priority is the *standby command switch* (SC). The other switches in the cluster standby group are the *passive command switches* (PC). If the active command switch and the standby command switch become disabled *at the same time*, the passive command switch with the highest priority becomes the active command switch. For the limitations to automatic discovery, see the "Automatic Recovery of Cluster Configuration" section on page 5-13. For information about changing HSRP priority values, see the "Configuring HSRP Priority" section on page 23-6. The commands are the same for changing the priority of cluster standby group members and router-redundancy group members.

Note The HSRP standby hold time interval should be greater than or equal to 3 times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and hello time intervals, see the "Configuring HSRP Authentication and Timers" section on page 23-8.

These connectivity guidelines ensure automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices. These topics also provide more detail about standby command switches:

- "Virtual IP Addresses" section on page 5-13
- "Automatic Recovery of Cluster Configuration" section on page 5-13
- "Considerations for Cluster Standby Groups" section on page 5-14

Virtual IP Addresses

You need to assign a unique virtual IP address and group number and name to the cluster standby group. This information must be configured on a specificVLAN or routed port on the active command switch. The active command switch receives traffic destined for the virtual IP address. To manage the cluster, you must access the active command switch through the virtual IP address, not through the command-switch IP address. This is in case the IP address of the active command switch is different from the virtual IP address of the cluster standby group.

If the active command switch fails, the standby command switch assumes ownership of the virtual IP address and becomes the active command switch. The passive switches in the cluster standby group compare their assigned priorities to determine the new standby command switch. The passive standby switch with the highest priority then becomes the standby command switch. When the previously active command switch becomes active again, it resumes its role as the active command switch, and the current active command switch becomes the standby command switch again. For more information about IP address in switch clusters, see the "IP Addresses" section on page 5-15.

Automatic Recovery of Cluster Configuration

The active command switch continually forwards cluster-configuration information (but not device-configuration information) to the standby command switch. This ensures that the standby command switch can take over the cluster immediately after the active command switch fails.

Automatic discovery has these limitations:

- This limitation applies only to clusters that have Catalyst 2950 and Catalyst 3550 command and standby command switches: If the active command switch and standby command switch become disabled *at the same time*, the passive command switch with the highest priority becomes the active command switch. However, because it was a passive standby command switch, the previous command switch *did not* forward cluster-configuration information to it. The active command switch only forwards cluster-configuration information to the standby command switch. You must therefore rebuild the cluster.
- This limitation applies to all clusters: If the active command switch fails and there are more than two switches in the cluster standby group, the new command switch does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL member switches. You must re-add these member switches to the cluster.
- This limitation applies to all clusters: If the active command switch fails and becomes active again, it does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL member switches. You must re-add these member switches to the cluster.

When the previously active command switch resumes its active role, it receives a copy of the latest cluster configuration from the active command switch, including members that were added while it was down. The active command switch sends a copy of the cluster configuration to the cluster standby group.

Considerations for Cluster Standby Groups

In addition to providing a virtual IP address to the cluster standby group, these requirements apply:

- When the command switch is a Catalyst 3550 switch, all standby command switches must be Catalyst 3550 switches or Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.

  When the command switch is a Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.

  When the Catalyst 2950 command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches also running Cisco IOS Release 12.0(5)WC2 or earlier: Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches.

  We recommend that the standby command switches be of the same platform family as the command switch. For example, a cluster with a Catalyst 3550 command switch should have Catalyst 3550 standby command switches.

- Only one cluster standby group can be assigned to a cluster. You can have more than one router-redundancy standby group.

  An HSRP group can be both a cluster standby group and a router-redundancy group. However, if a router-redundancy group becomes a cluster standby group, router redundancy becomes disabled on that group. You can reenable it by using the CLI. For more information about HSRP and router redundancy, see Chapter 23, "Configuring HSRP."

- All standby-group members must be members of the cluster.

  Note There is no limit to the number of switches you can assign as standby command switches. However, the total number of switches in the cluster—which would include the active command switch, standby-group members, and member switches—cannot be more than 16.

- The active command switch must be connected through the same VLAN to each standby-group member (Figure 5-8). Each standby-group member must also be redundantly connected to each other through at least one VLAN in common with the switch cluster.

  Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL member switches must be connected to the cluster standby group through their management VLANs.

  For more information about VLANs in switch clusters, see these sections:

  – "Discovery through Different VLANs" section on page 5-7
  – "Discovery through Different Management VLANs" section on page 5-9
  – "Discovery through the Same Management VLAN" section on page 5-8

Figure 5-8 VLAN Connectivity between Standby-Group Members and Cluster Members

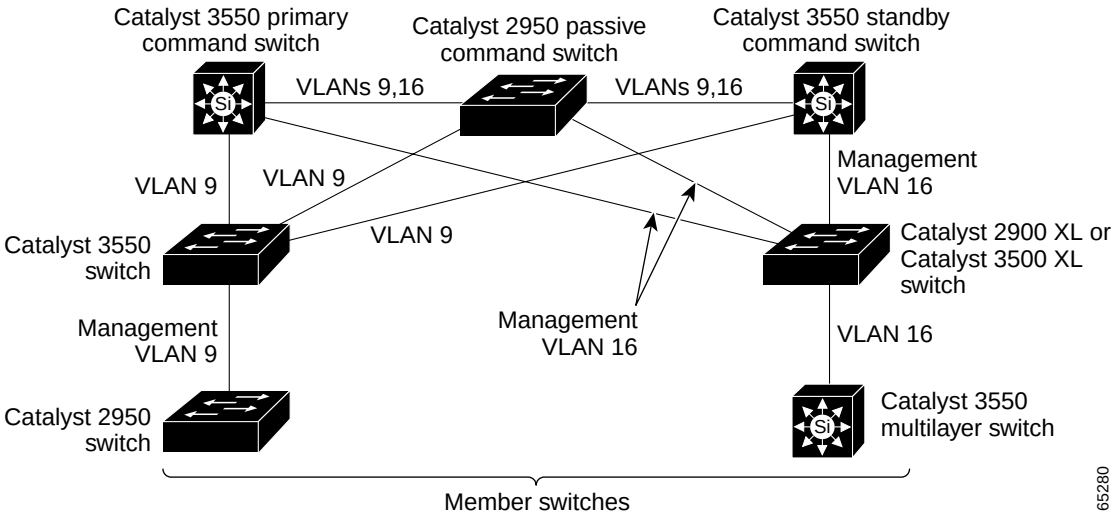

IP Addresses

You must assign IP information to a command switch. You can assign more than one IP address to the command switch, and you can access the cluster through any of the command-switch IP addresses. If you configure a cluster standby group, you must use the standby-group virtual IP address to manage the cluster from the active command switch. Using the virtual IP address ensures that you retain connectivity to the cluster if the active command switch fails and that a standby command switch becomes the active command switch.

If the active command switch fails and the standby command switch takes over, you must either use the standby-group virtual IP address to access the cluster or any of the IP addresses available on the new active command switch.

You can assign an IP address to a cluster-capable switch, but it is not necessary. A member switch is managed and communicates with other member switches through the command-switch IP address. If the member switch leaves the cluster and it does not have its own IP address, you then must assign IP information to it to manage it as a standalone switch.

Note Changing the command switch IP address ends your CMS session on the switch. Restart your CMS session by entering the new IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Internet Explorer), as described in the release notes.

For more information about IP addresses, see Chapter 4, "Assigning the Switch IP Address and Default Gateway."

Host Names

You do not need to assign a host name to either a command switch or an eligible cluster member. However, a host name assigned to the command switch can help to more easily identify the switch cluster. The default host name for the switch is *Switch*.

If a switch joins a cluster and it does not have a host name, the command switch appends a unique member number to its own host name and assigns it sequentially as each switch joins the cluster. The number means the order in which the switch was added to the cluster. For example, a command switch named *eng-cluster* could name the fifth cluster member *eng-cluster-5*.

If a switch has a host name, it retains that name when it joins a cluster. It retains that host name even after it leaves the cluster.

If a switch received its host name from the command switch, was removed from a cluster, was then added to a new cluster, and kept the same member number (such as *5*), the old host name (such as *eng-cluster-5*) is overwritten with the host name of the command switch in the new cluster (such as *mkg-cluster-5*). If the switch member number changes in the new cluster (such as *3*), the switch retains the previous name (*eng-cluster-5*).

Passwords

It is not necessary to assign passwords to an individual switch if it will be a cluster member. When a switch joins a cluster, it inherits the command-switch password and retains it when it leaves the cluster. If no command-switch password is configured, the member switch inherits a null password. Member switches only inherit the command-switch password.

If you change the member-switch password to be different from the command-switch password and save the change, the switch is not manageable by the command switch until you change the member-switch password to match the command-switch password. Rebooting the member switch does not revert the password back to the command-switch password. We recommend that you do not change the member-switch password after it joins a cluster.

For more information about passwords, see the "Preventing Unauthorized Access to Your Switch" section on page 6-1.

For password considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

SNMP Community Strings

A member switch inherits the command-switch first read-only (RO) and read-write (RW) community strings with *@esN* appended to the community strings:

- *command-switch-readonly-community-string@esN*, where *N* is the member-switch number.
- *command-switch-readwrite-community-string@esN*, where *N* is the member-switch number.

If the command switch has multiple read-only or read-write community strings, only the first read-only and read-write strings are propagated to the member switch.

The switches support an unlimited number of community strings and string lengths. For more information about SNMP and community strings, see Chapter 18, "Configuring SNMP."

For SNMP considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides specific for those switches.

TACACS+

If Terminal Access Controller Access Control System Plus (TACACS+) is configured on the command switch, TACACS+ must also be configured on all member switches to access the switch cluster from CMS. For more information about TACACS+, see the "Controlling Switch Access with TACACS+" section on page 6-10.

Access Modes in CMS

CMS provides two levels of access to the configuration options: read-write access and read-only access. Privilege levels 0 to 15 are supported.

- Privilege level 15 provides you with read-write access to CMS.
- Privilege levels 1 to 14 provide you with read-only access to CMS. Any options in the CMS windows, menu bar, toolbar, and popup menus that change the switch or cluster configuration are not shown in read-only mode.
- Privilege level 0 denies access to CMS.

For more information about CMS access modes, see the "Access Modes in CMS" section on page 3-31.

Note

- If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:
  - Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier
  - Catalyst 3550 member switches running Cisco IOS Release 12.1(6)EA1 or earlier

  For more information about this limitation, refer to the Catalyst 3550 release notes.
- These switches do not support read-only mode on CMS:
  - Catalyst 1900 and Catalyst 2820
  - Catalyst 2900 XL switches with 4-MB CPU DRAM

  In read-only mode, these switches appear as unavailable devices and cannot be configured from CMS.

LRE Profiles

A configuration conflict occurs if a switch cluster has LRE switches using both private and public profiles. If one LRE switch in a cluster is assigned a public profile, all LRE switches in that cluster must have that same public profile. Before you add an LRE switch to a cluster, make sure that you assign it the same public profile used by other LRE switches in the cluster.

A cluster can have a mix of LRE switches using different private profiles.

For more information about the Catalyst 2900 LRE XL switches and LRE technology, refer to the Catalyst 2900 XL and Catalyst 3500 XL documentation for Cisco IOS Release 12.0(5)WC2.

Availability of Switch-Specific Features in Switch Clusters

The menu bar on the command switch displays all options available from the switch cluster. Therefore, features specific to a member switch are available from the command-switch menu bar. For example, **Device** > **LRE Profile** appears in the command-switch menu bar when at least one Catalyst 2900 LRE XL switch is in the cluster.

Creating a Switch Cluster

Using CMS to create a cluster is easier than using the CLI commands. This section provides this information:

- "Enabling a Command Switch" section on page 5-19
- "Adding Member Switches" section on page 5-20
- "Creating a Cluster Standby Group" section on page 5-22
- "Verifying a Switch Cluster" section on page 5-24

This section assumes you have already cabled the switches, as described in the switch hardware installation guide, and followed the guidelines described in the "Planning a Switch Cluster" section on page 5-4.

Note Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be command switches and which ones can only be member switches, and for the required software versions and browser and Java plug-in configurations.

Enabling a Command Switch

The switch you designate to be the command switch must meet the requirements described in the "Command Switch Characteristics" section on page 5-2, "Planning a Switch Cluster" section on page 5-4, and the release notes.

Note

- We strongly recommend that the highest-end, command-capable switch in the cluster be the command switch:
- If your switch cluster has a Catalyst 3550 switch, that switch should be the command switch.
- If your switch cluster has Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches, the Catalyst 2950 should be the command switch.
- If your switch cluster has Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, and Catalyst 3500 XL switches, either the Catalyst 2900 XL or Catalyst 3500 XL should be the command switch.

You can enable a command switch, name the cluster, and assign an IP address and a password to the command switch when you run the setup program during initial switch setup. For information about using the setup program, refer to the release notes.

If you did not enable a command switch during initial switch setup, launch Device Manager from a command-capable switch, and select **Cluster > Create Cluster**. Enter a cluster number (the default is 0), and use up to 31 characters to name the cluster (Figure 5-9). Instead of using CMS to enable a command switch, you can use the **cluster enable** global configuration command.

Figure 5-9 Create Cluster Window

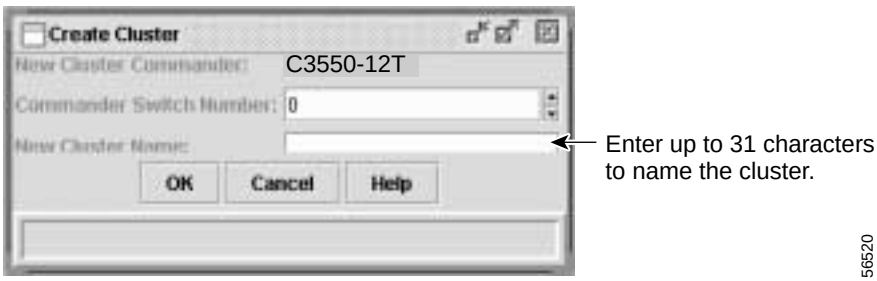

Adding Member Switches

As explained in the "Automatic Discovery of Cluster Candidates and Members" section on page 5-4, the command switch automatically discovers candidate switches. When you add new cluster-capable switches to the network, the command switch discovers and adds them to a list of candidate switches. To display an updated cluster candidates list from the Add to Cluster window (Figure 5-10), either relaunch CMS and redisplay this window, or follow these steps:

1. Close the Add to Cluster window.
2. Select **View** > **Refresh**.
3. Select **Cluster** > **Add to Cluster** to redisplay the Add to Cluster window.

From CMS, there are two ways to add switches to a cluster:

- Select **Cluster** > **Add to Cluster**, select a candidate switch from the list, click **Add**, and click **OK**. To add more than one candidate switch, press **Ctrl**, and make your choices, or press **Shift**, and choose the first and last switch in a range.
- Display the Topology view, right-click a candidate-switch icon, and select **Add to Cluster** (Figure 5-11). In the Topology view, candidate switches are cyan, and member switches are green. To add more than one candidate switch, press **Ctrl**, and left-click the candidates that you want to add.

Instead of using CMS to add members to the cluster, you can use the **cluster setup** privileged EXEC command or **cluster member** global configuration command from the command switch.

You can select one or more switches as long as the total number of switches in the cluster does not exceed 16 (this includes the command switch). When a cluster has 16 members, the **Add to Cluster** option is not available for that cluster. In this case, you must remove a member switch before adding a new one.

If a password has been configured on a candidate switch, you are prompted to enter it before you can add it to the cluster. If the candidate switch does not have a password, any entry is ignored.

If multiple candidates switches have the same password, you can select them as a group, and add them at the same time. If you incorrectly enter the password for those candidates, none of them is added to the cluster. If a candidate switch in the group has a password different from the group, only that specific candidate switch is not added to the cluster.

When a candidate switch joins a cluster, it inherits the command-switch password. For more information about setting passwords, see the "Passwords" section on page 5-16.

Figure 5-10 Add to Cluster Window

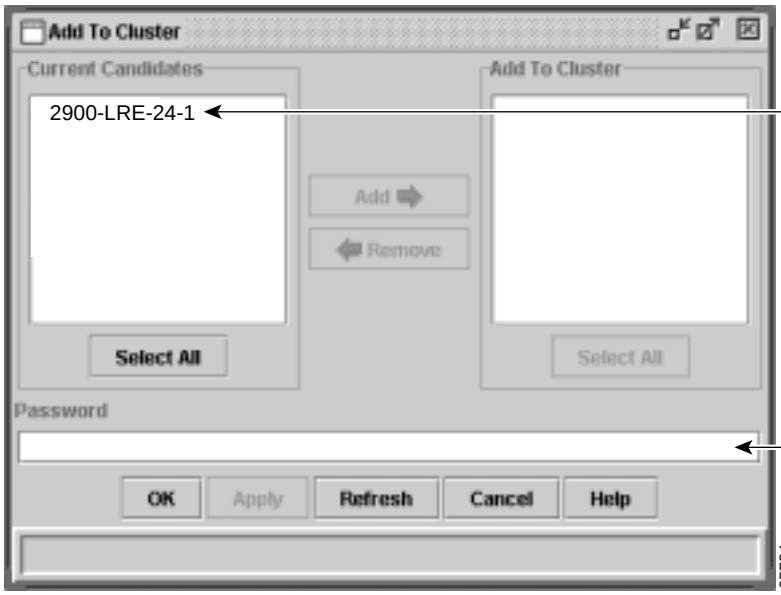

Select a switch, and click **Add**. Press Ctrl and left-click to select more than one switch.

Enter the password of the candidate switch. If no password exists for the switch, leave this field blank.

Figure 5-11 Using the Topology View to Add Member Switches

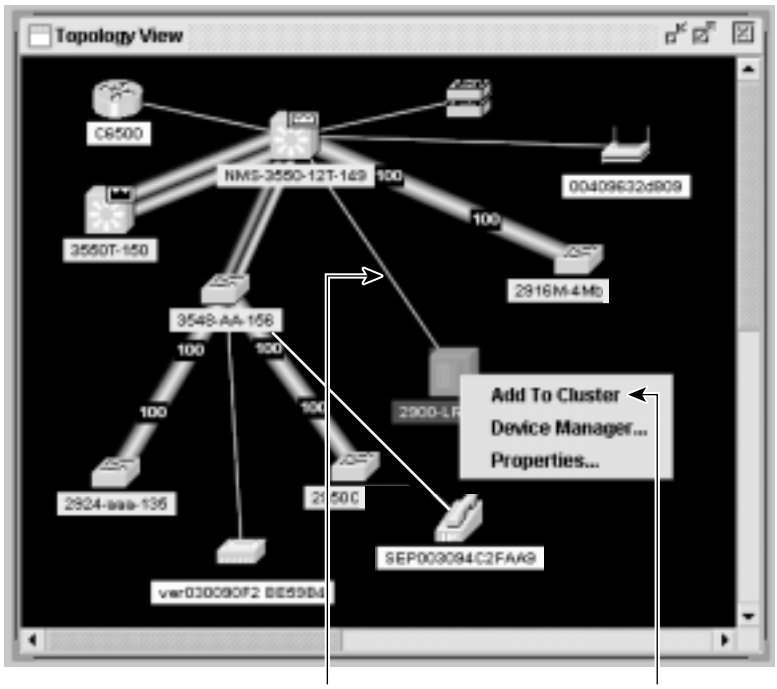

Thin line means a connection to a candidate switch.

Right-click a candidate switch to display the pop-up menu, and select **Add to Cluster** to add the switch to the cluster.

Creating a Cluster Standby Group

The cluster standby group members must meet the requirements described in the "Standby Command Switch Characteristics" section on page 5-3 and "HSRP and Standby Command Switches" section on page 5-12. To create a cluster standby group, select **Cluster** > **Standby Command Switches** (Figure 5-12).

Instead of using CMS to add switches to a standby group and to bind the standby group to a cluster, you can use the **standby ip**, **standby name**, **standby priority** interface configuration commands and the **cluster standby group** global configuration command.

Note

- When the command switch is a Catalyst 3550 switch, all standby command switches must be Catalyst 3550 switches or Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the command switch is a Catalyst 2950 switch running Cisco IOS Release 12.1(6)EA2 or later, all standby command switches must be Catalyst 2950 switches running Cisco IOS Release 12.1(6)EA2 or later.
- When the Catalyst 2950 command switch is running Cisco IOS Release 12.0(5)WC2 or earlier, the standby command switches can be these switches also running Cisco IOS Release 12.0(5)WC2 or earlier: Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches.

These abbreviations are appended to the switch host names in the Standby Command Group list to show their eligibility or status in the cluster standby group:

- AC—Active command switch
- SC—Standby command switch
- PC—Member of the cluster standby group but not the standby command switch
- HC—Candidate switch that can be added to the cluster standby group
- CC—Command switch when HSRP is disabled

You must enter a virtual IP address for the cluster standby group. This address must be in the same subnet as the IP addresses of the switch. The group number must be unique within the IP subnet. It can be from 0 to 255, and the default is 0. The group name can have up to 31 characters.

The Standby Command Configuration window uses the default values for the **preempt** and **name** commands that you have set by using the CLI. If you use this window to create the HSRP group, all switches in the group have the **preempt** command enabled. You must also provide a name for the group.

Note The HSRP standby hold time interval should be greater than or equal to 3 times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and hello time intervals, see the "Configuring HSRP Authentication and Timers" section on page 23-8.

Figure 5-12 Standby Command Configuration Window

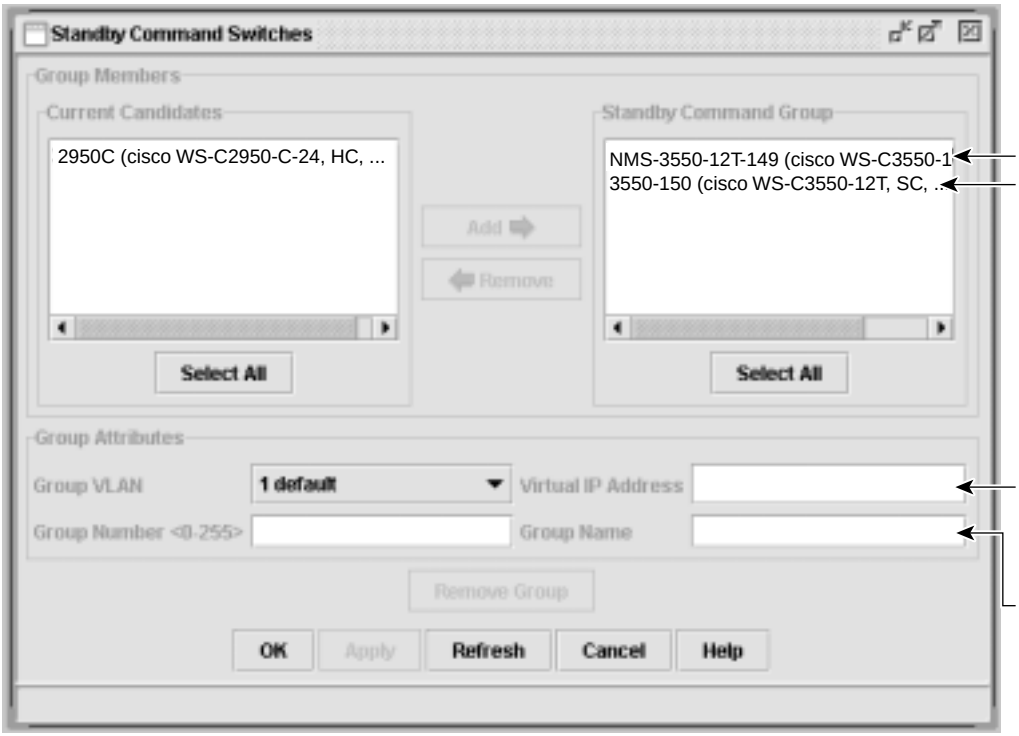

Verifying a Switch Cluster

When you finish adding cluster members, follow these steps to verify the cluster:

Step 1 Enter the command switch IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Microsoft Internet Explorer) to access all switches in the cluster.

Step 2 Enter the command-switch password.

Step 3 Select **View** > **Topology** to display the topology of the switch cluster and to view link information (Figure 3-6 on page 3-11). For complete information about the Topology view, including descriptions of the icons, links, and colors used in the Topology view, see the "Topology View" section on page 3-10.

Step 4 Select **Reports** > **Inventory** to display an inventory of the switches in the cluster (Figure 5-13).

The summary includes information such as switch model numbers, serial numbers, software versions, IP information, and location.

You can also display port and switch statistics from **Reports** > **Port Statistics** and **Port** > **Port Settings** > **Runtime Status**.

Instead of using CMS to verify the cluster, you can use the **show cluster members** user EXEC command from the command switch or use the **show cluster** user EXEC command from the command switch or from a member switch.

Figure 5-13 Inventory Window

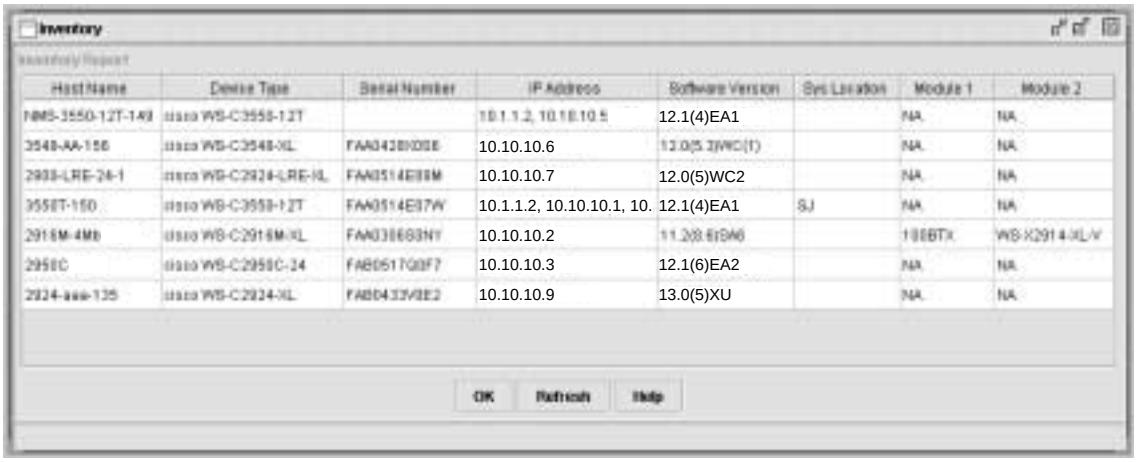

If you lose connectivity with a member switch or if a command switch fails, see the "Using Recovery Procedures" section on page 27-1.

For more information about creating and managing clusters, refer to the online help. For information about the cluster commands, refer to the switch command reference.

Using the CLI to Manage Switch Clusters

You can configure member switches from the CLI by first logging into the command switch. Enter the **rcommand** user EXEC command and the member switch number to start a Telnet session (through a console or Telnet connection) and to access the member switch CLI. After this, the command mode changes, and the IOS commands operate as usual. Enter the **exit** privileged EXEC command on the member switch to return to the command-switch CLI.

This example shows how to log into member-switch 3 from the command-switch CLI:

```
switch# rcommand 3
```

If you do not know the member-switch number, enter the **show cluster members** privileged EXEC command on the command switch. For more information about the **rcommand** command and all other cluster commands, refer to the switch command reference.

The Telnet session accesses the member-switch CLI at the same privilege level as on the command switch. The IOS commands then operate as usual. For instructions on configuring the switch for a Telnet session, see the "Setting a Telnet Password for a Terminal Line" section on page 6-6.

Catalyst 1900 and Catalyst 2820 CLI Considerations

If your switch cluster has Catalyst 1900 and Catalyst 2820 switches running standard edition software, the Telnet session accesses the management console (a menu-driven interface) if the command switch is at privilege level 15. If the command switch is at privilege level 1 to 14, you are prompted for the password to access the menu console.

Command-switch privilege levels map to the Catalyst 1900 and Catalyst 2820 member switches running standard and Enterprise Edition Software as follows:

- If the command-switch privilege level is 1 to 14, the member switch is accessed at privilege level 1.
- If the command-switch privilege level is 15, the member switch is accessed at privilege level 15.

Note The Catalyst 1900 and Catalyst 2820 CLI is available only on switches running Enterprise Edition Software.

For more information about the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

Using SNMP to Manage Switch Clusters

When you first power on the switch, SNMP is enabled if you enter the IP information by using the setup program and accept its proposed configuration. If you did not use the setup program to enter the IP information and SNMP was not enabled, you can enable it as described in the "Configuring SNMP" section on page 18-4. On Catalyst 1900 and Catalyst 2820 switches, SNMP is enabled by default.

When you create a cluster, the command switch manages the exchange of messages between member switches and an SNMP application. The cluster software on the command switch appends the member switch number (*@esN*, where *N* is the switch number) to the first configured read-write and read-only community strings on the command switch and propagates them to the member switch. The command switch uses this community string to control the forwarding of gets, sets, and get-next messages between the SNMP management station and the member switches.

Note When a cluster standby group is configured, the command switch can change without your knowledge. Use the first read-write and read-only community strings to communicate with the command switch if there is a cluster standby group configured for the cluster.

If the member switch does not have an IP address, the command switch redirects traps from the member switch to the management station, as shown in Figure 5-14. If a member switch has its own IP address and community strings, the member switch can send traps directly to the management station, without going through the command switch.

If a member switch has its own IP address and community strings, they can be used in addition to the access provided by the command switch. For more information about SNMP and community strings, see Chapter 18, "Configuring SNMP."

Figure 5-14 SNMP Management for a Cluster

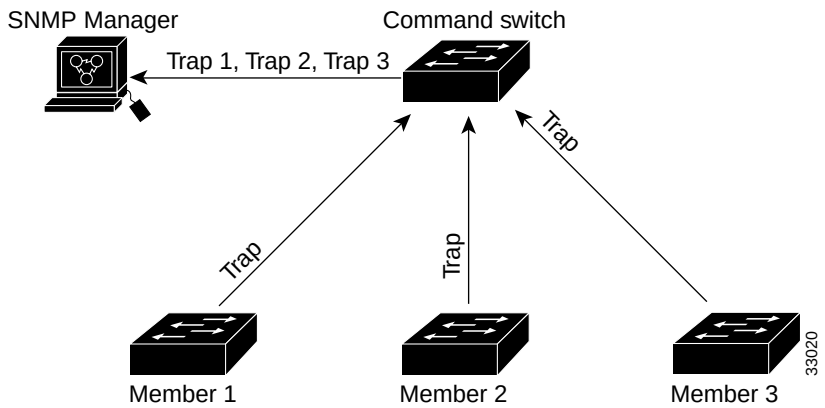

Administering the Switch

This chapter describes how to perform one-time operations to administer your switch. This chapter consists of these sections:

- Preventing Unauthorized Access to Your Switch, page 6-1
- Protecting Access to Privileged EXEC Commands, page 6-2
- Controlling Switch Access with TACACS+, page 6-10
- Controlling Switch Access with RADIUS, page 6-17
- Configuring the Switch for Local Authentication and Authorization, page 6-31
- Managing the System Time and Date, page 6-32
- Configuring a System Name and Prompt, page 6-46
- Creating a Banner, page 6-49
- Managing the MAC Address Table, page 6-51
- Optimizing System Resources for User-Selected Features, page 6-57

Preventing Unauthorized Access to Your Switch

You can prevent unauthorized users from reconfiguring your switch and viewing configuration information. Typically, you want network administrators to have access to your switch while you restrict access to users who dial from outside the network through an asynchronous port, connect from outside the network through a serial port, or connect through a terminal or workstation from within the local network.

To prevent unauthorized access into your switch, you should configure one or more of these security features:

- At a minimum, you should configure passwords and privileges at each switch port. These passwords are locally stored on the switch. When users attempt to access the switch through a port or line, they must enter the password specified for the port or line before they can access the switch. For more information, see the "Protecting Access to Privileged EXEC Commands" section on page 6-2.
- For an additional layer of security, you can also configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair. For more information, see the "Configuring Username and Password Pairs" section on page 6-7.
- If you want to use username and password pairs, but you want to store them centrally on a server instead of locally, you can store them in a database on a security server. Multiple networking devices can then use the same database to obtain user authentication (and, if necessary, authorization) information. For more information, see the "Controlling Switch Access with TACACS+" section on page 6-10.

Protecting Access to Privileged EXEC Commands

A simple way of providing terminal access control in your network is to use passwords and assign privilege levels. Password protection restricts access to a network or network device. Privilege levels define what commands users can issue after they have logged into a network device.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section describes how to control access to the configuration file and privileged EXEC commands. It contains this configuration information:

- Default Password and Privilege Level Configuration, page 6-3
- Setting or Changing a Static Enable Password, page 6-3
- Protecting Enable and Enable Secret Passwords with Encryption, page 6-4
- Setting a Telnet Password for a Terminal Line, page 6-6
- Configuring Username and Password Pairs, page 6-7
- Configuring Multiple Privilege Levels, page 6-8

Default Password and Privilege Level Configuration

Table 6-1 shows the default password and privilege level configuration.

Table 6-1 Default Password and Privilege Levels

Feature	Default Setting
Enable password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is not encrypted in the configuration file.
Enable secret password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is encrypted before it is written to the configuration file.
Line password	No password is defined.

Setting or Changing a Static Enable Password

The enable password controls access to the privileged EXEC mode. Beginning in privileged EXEC mode, follow these steps to set or change a static enable password:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**enable password** *password*	Define a new password or change an existing password for access to privileged EXEC mode. By default, no password is defined. For *password*, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. It can contain the question mark (?) character if you precede the question mark with the key combination Crtl-v when you create the password; for example, to create the password abc?123, do this: Enter **abc**. Enter **Crtl-v**. Enter **?123**. When the system prompts you to enter the enable password, you need not precede the question mark with the Ctrl-v; you can simply enter abc?123 at the password prompt.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file. The enable password is not encrypted and can be read in the switch configuration file.

To remove the password, use the **no enable password** global configuration command.

This example shows how to change the enable password to *l1u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

```
Switch(config)# enable password l1u2c3k4y5
```

Protecting Enable and Enable Secret Passwords with Encryption

To provide an additional layer of security, particularly for passwords that cross the network or that are stored on a Trivial File Transfer Protocol (TFTP) server, you can use either the **enable password** or **enable secret** global configuration commands. Both commands accomplish the same thing; that is, you can establish an encrypted password that users must enter to access privileged EXEC mode (the default) or any privilege level you specify.

We recommend that you use the **enable secret** command because it uses an improved encryption algorithm.

If you configure the **enable secret** command, it takes precedence over the **enable password** command; the two commands cannot be in effect simultaneously.

Beginning in privileged EXEC mode, follow these steps to configure encryption for enable and enable secret passwords:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**enable password** [**level** *level*] {*password* | *encryption-type encrypted-password*}	Define a new password or change an existing password for access to privileged EXEC mode.
	or	or
	enable secret [**level** *level*] {*password* | *encryption-type encrypted-password*}	Define a secret password, which is saved using a nonreversible encryption method. • (Optional) For *level*, the range is from 0 to 15. Level 1 is normal user EXEC mode privileges. The default level is 15 (privileged EXEC mode privileges). • For *password*, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined. • (Optional) For *encryption-type*, only type 5, a Cisco proprietary encryption algorithm, is available. If you specify an encryption type, you must provide an encrypted password—an encrypted password you copy from another Catalyst 3550 multilayer switch configuration. **Note** If you specify an encryption type and then enter a clear text password, you can not re-enter privileged EXEC mode. You cannot recover a lost encrypted password by any method.

	Command	Purpose
Step 3	**service password-encryption**	(Optional) Encrypt the password when the password is defined or when the current configuration is written. Encryption prevents the password from being readable in the configuration file.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

If both the enable and enable secret passwords are defined, users must enter the enable secret password.

Use the **level** keyword to define a password for a specific privilege level. After you specify the level and set a password, give the password only to users who need to have access at this level. Use the **privilege level** global configuration command to specify commands accessible at various levels. For more information, see the "Configuring Multiple Privilege Levels" section on page 6-8.

If you enable password encryption, it applies to all passwords including username passwords, authentication key passwords, the privileged command password, and console and virtual terminal line passwords.

To remove a password and level, use the **no enable password** [**level** *level*] or **no enable secret** [**level** *level*] global configuration command. To disable password encryption, use the **no service password-encryption** global configuration command.

This example shows how to configure the encrypted password *1FaD0$Xyti5Rkls3LoyxzS8* for privilege level 2:

```
Switch(config)# enable secret level 2 5 $1$FaD0$Xyti5Rkls3LoyxzS8
```

Disabling Password Recovery

The default configuration for Catalyst 3550 switches allows an end user with physical access to the switch to recover from a lost password by interrupting the boot process while the switch is powering up and then by entering a new password. The password recovery disable feature for Catalyst 3550 Fast Ethernet switches allows the system administrator to protect access to the switch password by disabling part of this functionality and allowing the user to interrupt the boot process only by agreeing to set the system back to the default configuration. With password recovery disabled, you can still interrupt the boot process and change the password, but the configuration file (config.text) and the VLAN database file (vlan.dat) are deleted.

Note The password recovery disable feature is valid only on Catalyst 3550 Fast Ethernet switches; it is not available for Catalyst 3550 Gigabit Ethernet switches.

Note If you disable password recovery, we recommend that you keep a backup copy of the configuration file on a secure server in case the end user interrupts the boot process and sets the system back to defaults. Do not keep a backup copy of the configuration file on the switch. If the switch is operating in VTP transparent mode, we recommend that you also keep a backup copy of the VLAN database file on a secure server. When the switch is returned to the default system configuration, you can download the saved files to the switch by using the XMODEM protocol. For more information, see the "Recovering from a Lost or Forgotten Password" section on page 27-3.

Beginning in privileged EXEC mode, follow these steps to disable password recovery:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no service password-recovery**	Disable password recovery. This setting is saved in an area of the Flash memory that is accessible by the boot loader and the IOS image, but it is not part of the file system and is not accessible by any user.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show version**	Verify the configuration by checking the last few lines of the display.

To re-enable password recovery, use the **service password-recovery** global configuration command.

Note Disabling password recovery will not work if you have set the switch to boot manually by using the **boot manual** global configuration command because this command allows the user to automatically see the boot loader prompt (*switch:*) after power cycling the switch.

Setting a Telnet Password for a Terminal Line

When you power-up your switch for the first time, an automatic setup program runs to assign IP information and to create a default configuration for continued use. The setup program also prompts you to configure your switch for Telnet access through a password. If you neglected to configure this password during the setup program, you can configure it now through the command-line interface (CLI).

Beginning in privileged EXEC mode, follow these steps to configure your switch for Telnet access:

	Command	Purpose
Step 1		Attach a PC or workstation with emulation software to the switch console port. The default data characteristics of the console port are 9600, 8, 1, no parity. You might need to press the Return key several times to see the command-line prompt.
Step 2	**enable password** *password*	Enter privileged EXEC mode.
Step 3	**configure terminal**	Enter global configuration mode.
Step 4	**line vty 0 15**	Configure the number of Telnet sessions (lines), and enter line configuration mode. There are 16 possible sessions on a command-capable switch. The 0 and 15 mean that you are configuring all 16 possible Telnet sessions.
Step 5	**password** *password*	Enter a Telnet password for the line or lines. For *password*, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 6	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 7	**show running-config**	Verify your entries. The password is listed under the command **line vty 0 15**.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the password, use the **no password** global configuration command.

This example shows how to set the Telnet password to *let45me67in89*:

```
Switch(config)# line vty 10
Switch(config-line)# password let45me67in89
```

Configuring Username and Password Pairs

You can configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair.

Beginning in privileged EXEC mode, follow these steps to establish a username-based authentication system that requests a username and password during switch login operations:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**username** *name* [**privilege** *level*] {**password** *encryption-type password*}	Enter the username, privilege level, and password for each user. • For *name*, specify the user ID as one word. Spaces and quotation marks are not allowed. • (Optional) For *level*, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 1 gives user EXEC mode access. • For *encryption-type*, enter 0 to specify that an unencrypted password will follow. Enter 7 to specify that a hidden password will follow. • For *password*, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the **username** command.
Step 3	**line console 0** or **line vty 0 15**	Enter line configuration mode, and configure the console port (line 0) or the VTY lines (line 0 to 15).
Step 4	**login local**	Enable local password checking at login time. Authentication is based on the username specified in Step 2.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable username authentication for a specific user, use the **no username** *name* global configuration command. To disable password checking and allow connections without a password, use the **no login** line configuration command.

Configuring Multiple Privilege Levels

By default, the IOS software has two modes of password security: user EXEC and privileged EXEC. You can configure up to 16 hierarchical levels of commands for each mode. By configuring multiple passwords, you can allow different sets of users to have access to specified commands.

For example, if you want many users to have access to the **clear line** command, you can assign it level 2 security and distribute the level 2 password fairly widely. But if you want more restricted access to the **configure** command, you can assign it level 3 security and distribute that password to a more restricted group of users.

This section includes this configuration information:

- Setting the Privilege Level for a Command, page 6-8
- Changing the Default Privilege Level for Lines, page 6-9
- Logging into and Exiting a Privilege Level, page 6-10

Setting the Privilege Level for a Command

Beginning in privileged EXEC mode, follow these steps to set the privilege level for a command mode:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**privilege** *mode* **level** *level command*	Set the privilege level for a command. • For *mode*, enter **configure** for global configuration mode, **exec** for EXEC mode, **interface** for interface configuration mode, or **line** for line configuration mode. • For *level*, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the **enable** password. • For *command*, specify the command to which you want to restrict access.
Step 3	**enable password level** *level password*	Specify the enable password for the privilege level. • For *level*, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. • For *password*, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show running-config** or **show privilege**	Verify your entries. The first command displays the password and access level configuration. The second command displays the privilege level configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

When you set a command to a privilege level, all commands whose syntax is a subset of that command are also set to that level. For example, if you set the **show ip route** command to level 15, the **show** commands and **show ip** commands are automatically set to privilege level 15 unless you set them individually to different levels.

To return to the default privilege for a given command, use the **no privilege** *mode* **level** *level command* global configuration command.

This example shows how to set the **configure** command to privilege level 14 and define *SecretPswd14* as the password users must enter to use level 14 commands:

```
Switch(config)# privilege exec level 14 configure
Switch(config)# enable password level 14 SecretPswd14
```

Changing the Default Privilege Level for Lines

Beginning in privileged EXEC mode, follow these steps to change the default privilege level for a line:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**line vty** *line*	Select the virtual terminal line on which to restrict access.
Step 3	**privilege level** *level*	Change the default privilege level for the line. For *level*, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the **enable** password.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config** or **show privilege**	Verify your entries. The first command displays the password and access level configuration. The second command displays the privilege level configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Users can override the privilege level you set using the **privilege level** line configuration command by logging in to the line and enabling a different privilege level. They can lower the privilege level by using the **disable** command. If users know the password to a higher privilege level, they can use that password to enable the higher privilege level. You might specify a high level or privilege level for your console line to restrict line usage.

To return to the default line privilege level, use the **no privilege level** line configuration command.

Logging into and Exiting a Privilege Level

Beginning in privileged EXEC mode, follow these steps to log in to a specified privilege level and to exit to a specified privilege level:

	Command	Purpose
Step 1	**enable** *level*	Log in to a specified privilege level. For *level*, the range is 0 to 15.
Step 2	**disable** *level*	Exit to a specified privilege level. For *level*, the range is 0 to 15.

Controlling Switch Access with TACACS+

This section describes how to enable and configure Terminal Access Controller Access Control System Plus (TACACS+), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. TACACS+ is facilitated through authentication, authorization, accounting (AAA) and can be enabled only through AAA commands.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section contains this configuration information:

- Understanding TACACS+, page 6-10
- TACACS+ Operation, page 6-12
- Configuring TACACS+, page 6-13
- Displaying the TACACS+ Configuration, page 6-17

Understanding TACACS+

TACACS+ is a security application that provides centralized validation of users attempting to gain access to your switch. TACACS+ services are maintained in a database on a TACACS+ daemon typically running on a UNIX or Windows NT workstation. You should have access to and should configure a TACACS+ server before the configuring TACACS+ features on your switch.

TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service—authentication, authorization, and accounting—independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The goal of TACACS+ is to provide a method for managing multiple network access points from a single management service. Your switch can be a network access server along with other Cisco routers and access servers. A network access server provides connections to a single user, to a network or subnetwork, and to interconnected networks as shown in Figure 6-1.

Figure 6-1 Typical TACACS+ Network Configuration

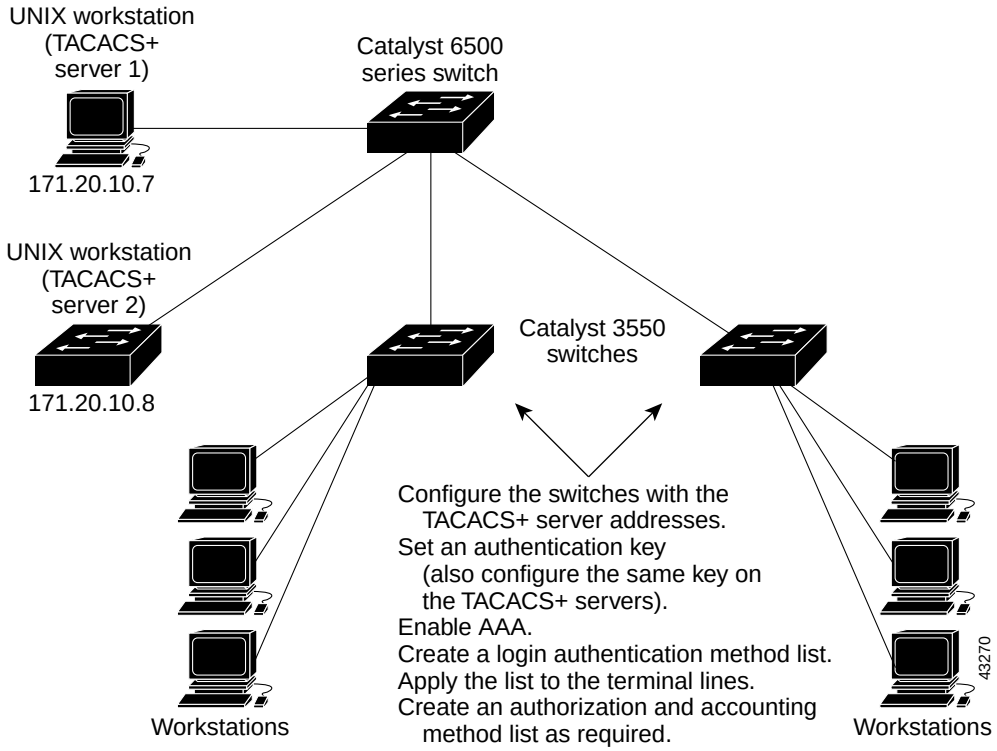

TACACS+, administered through the AAA security services, can provide these services:

- Authentication—Provides complete control of authentication through login and password dialog, challenge and response, and messaging support.

  The authentication facility can conduct a dialog with the user (for example, after a username and password are provided, to challenge a user with several questions, such as home address, mother's maiden name, service type, and social security number). The TACACS+ authentication service can also send messages to user screens. For example, a message could notify users that their passwords must be changed because of the company's password aging policy.

- Authorization—Provides fine-grained control over user capabilities for the duration of the user's session, including but not limited to setting autocommands, access control, session duration, or protocol support. You can also enforce restrictions on what commands a user can execute with the TACACS+ authorization feature.

- Accounting—Collects and sends information used for billing, auditing, and reporting to the TACACS+ daemon. Network managers can use the accounting facility to track user activity for a security audit or to provide information for user billing. Accounting records include user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

The TACACS+ protocol provides authentication between the switch and the TACACS+ daemon, and it ensures confidentiality because all protocol exchanges between the switch and the TACACS+ daemon are encrypted.

You need a system running the TACACS+ daemon software to use TACACS+ on your switch.

TACACS+ Operation

When a user attempts a simple ASCII login by authenticating to a switch using TACACS+, this process occurs:

1. When the connection is established, the switch contacts the TACACS+ daemon to obtain a username prompt, which is then displayed to the user. The user enters a username, and the switch then contacts the TACACS+ daemon to obtain a password prompt. The switch displays the password prompt to the user, the user enters a password, and the password is then sent to the TACACS+ daemon.

   TACACS+ allows a conversation to be held between the daemon and the user until the daemon receives enough information to authenticate the user. The daemon prompts for a username and password combination, but can include other items, such as the user's mother's maiden name.

2. The switch eventually receives one of these responses from the TACACS+ daemon:
   - ACCEPT—The user is authenticated and service can begin. If the switch is configured to require authorization, authorization begins at this time.
   - REJECT—The user is not authenticated. The user can be denied access or is prompted to retry the login sequence, depending on the TACACS+ daemon.
   - ERROR—An error occurred at some time during authentication with the daemon or in the network connection between the daemon and the switch. If an ERROR response is received, the switch typically tries to use an alternative method for authenticating the user.
   - CONTINUE—The user is prompted for additional authentication information.

   After authentication, the user undergoes an additional authorization phase if authorization has been enabled on the switch. Users must first successfully complete TACACS+ authentication before proceeding to TACACS+ authorization.

3. If TACACS+ authorization is required, the TACACS+ daemon is again contacted, and it returns an ACCEPT or REJECT authorization response. If an ACCEPT response is returned, the response contains data in the form of attributes that direct the EXEC or NETWORK session for that user, determining the services that the user can access:
   - Telnet, rlogin, or privileged EXEC services
   - Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring TACACS+

This section describes how to configure your switch to support TACACS+. At a minimum, you must identify the host or hosts maintaining the TACACS+ daemon and define the method lists for TACACS+ authentication. You can optionally define method lists for TACACS+ authorization and accounting. A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used, thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

This section contains this configuration information:

- Default TACACS+ Configuration, page 6-13
- Identifying the TACACS+ Server Host and Setting the Authentication Key, page 6-13
- Configuring TACACS+ Login Authentication, page 6-14
- Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services, page 6-16
- Starting TACACS+ Accounting, page 6-17

Default TACACS+ Configuration

TACACS+ and AAA are disabled by default.

To prevent a lapse in security, you cannot configure TACACS+ through a network management application.When enabled, TACACS+ can authenticate users accessing the switch through the CLI.

Note Although TACACS+ configuration is performed through the CLI, the TACACS+ server authenticates HTTP connections that have been configured with a privilege level of 15.

Identifying the TACACS+ Server Host and Setting the Authentication Key

You can configure the switch to use a single server or AAA server groups to group existing server hosts for authentication. You can group servers to select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list and contains the list of IP addresses of the selected server hosts.

Beginning in privileged EXEC mode, follow these steps to identify the IP host or host maintaining TACACS+ server and optionally set the encryption key:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**tacacs-server host** *hostname* [**port** *integer*] [**timeout** *integer*] [**key** *string*]	Identify the IP host or hosts maintaining a TACACS+ server. Enter this command multiple times to create a list of preferred hosts. The software searches for hosts in the order in which you specify them. • For *hostname*, specify the name or IP address of the host. • (Optional) For **port** *integer*, specify a server port number. The default is port 49. The range is 1 to 65535. • (Optional) For **timeout** *integer*, specify a time in seconds the switch waits for a response from the daemon before it times out and declares an error. The default is 5 seconds. The range is 1 to 1000 seconds. • (Optional) For **key** *string*, specify the encryption key for encrypting and decrypting all traffic between the switch and the TACACS+ daemon. You must configure the same key on the TACACS+ daemon for encryption to be successful.
Step 3	**aaa new-model**	Enable AAA.
Step 4	**aaa group server tacacs+** *group-name*	(Optional) Define the AAA server-group with a group name. This command puts the switch in a server group subconfiguration mode.
Step 5	**server** *ip-address*	(Optional) Associate a particular TACACS+ server with the defined server group. Repeat this step for each TACACS+ server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show tacacs**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the specified TACACS+ server name or address, use the **no tacacs-server host** *hostname* global configuration command. To remove a server group from the configuration list, use the **no aaa group server tacacs+** *group-name* global configuration command. To remove the IP address of a TACACS+ server, use the **no server ip-address** server group subconfiguration command.

Configuring TACACS+ Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined. A defined method list overrides the default method list.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa new-model**	Enable AAA.
Step 3	**aaa authentication login** {**default** | *list-name*} *method1* [*method2...*]	Create a login authentication method list. • To create a default list that is used when a named list is *not* specified in the **login authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. • For *list-name*, specify a character string to name the list you are creating. • For *method1...*, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. Select one of these methods: • **line**—Use the line password for authentication. You must define a line password before you can use this authentication method. Use the **password** *password* line configuration command. • **local**—Use the local username database for authentication. You must enter username information into the database. Use the **username** *password* global configuration command. • **tacacs+**—Uses TACACS+ authentication. You must configure the TACACS+ server before you can use this authentication method.
Step 4	**line** [**console** | **tty** | **vty**] *line-number* [*ending-line-number*]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.
Step 5	**login authentication** {**default** | *list-name*}	Apply the authentication list to a line or set of lines. • If you specify **default**, use the default list created with the **aaa authentication login** command. • For *list-name*, specify the list created with the **aaa authentication login** command.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** {**default** | *list-name*} *method1* [*method2...*] global configuration command. To either disable TACACS+ authentication for logins or to return to the default value, use the **no login authentication** {**default** | *list-name*} line configuration command.

Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is located either in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **tacacs+** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec tacacs+ local** command sets these authorization parameters:

- Use TACACS+ for privileged EXEC access authorization if authentication was performed by using TACACS+.
- Use the local database if authentication was not performed by using TACACS+.

Note Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify TACACS+ authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa authorization network tacacs+**	Configure the switch for user TACACS+ authorization for all network-related service requests.
Step 3	**aaa authorization exec tacacs+**	Configure the switch for user TACACS+ authorization to determine if the user has privileged EXEC access. The **exec** keyword might return user profile information (such as **autocommand** information).
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization** {**network** | **exec**} *method1* global configuration command.

Starting TACACS+ Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable TACACS+ accounting for each Cisco IOS privilege level and for network services:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa accounting network start-stop tacacs+**	Enable TACACS+ accounting for all network-related service requests.
Step 3	**aaa accounting exec start-stop tacacs+**	Enable TACACS+ accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting** {**network** | **exec**} {**start-stop**} *method1...* global configuration command.

Displaying the TACACS+ Configuration

To display TACACS+ server statistics, use the **show tacacs** privileged EXEC command.

Controlling Switch Access with RADIUS

This section describes how to enable and configure the Remote Authentication Dial-In User Service (RADIUS), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. RADIUS is facilitated through AAA and can be enabled only through AAA commands.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section contains this configuration information:

- Understanding RADIUS, page 6-18
- RADIUS Operation, page 6-19
- Configuring RADIUS, page 6-19
- Displaying the RADIUS Configuration, page 6-30

Understanding RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco routers and switches (including Catalyst 3550 multilayer switches and Catalyst 2950 series switches) and send authentication requests to a central RADIUS server, which contains all user authentication and network service access information. The RADIUS host is normally a multiuser system running RADIUS server software from Cisco (Cisco Secure Access Control Server version 3.0), Livingston, Merit, Microsoft, or another software provider. For more information, refer to the RADIUS server documentation.

Use RADIUS in these network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS. For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Turnkey network security environments in which applications support the RADIUS protocol, such as in an access environment that uses a *smart card* access control system. In one case, RADIUS has been used with Enigma's security cards to validates users and to grant access to network resources.
- Networks already using RADIUS. You can add a Cisco switch containing a RADIUS client to the network. This might be the first step when you make a transition to a TACACS+ server.
- Network in which the user must only access a single service. Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to the network through a protocol such as IEEE 802.1X. For more information about this protocol, see Chapter 7, "Configuring 802.1X Port-Based Authentication."
- Networks that require resource accounting. You can use RADIUS accounting independently of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, showing the amount of resources (such as time, packets, bytes, and so forth) used during the session. An Internet service provider might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in these network security situations:

- Multiprotocol access environments. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Frame Control Protocol (NBFCP), NetWare Asynchronous Services Interface (NASI), or X.25 PAD connections.
- Switch-to-switch or router-to-router situations. RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one device to a non-Cisco device if the non-Cisco device requires authentication.
- Networks using a variety of services. RADIUS generally binds a user to one service model.

Figure 6-2 Typical AAA Network Configuration

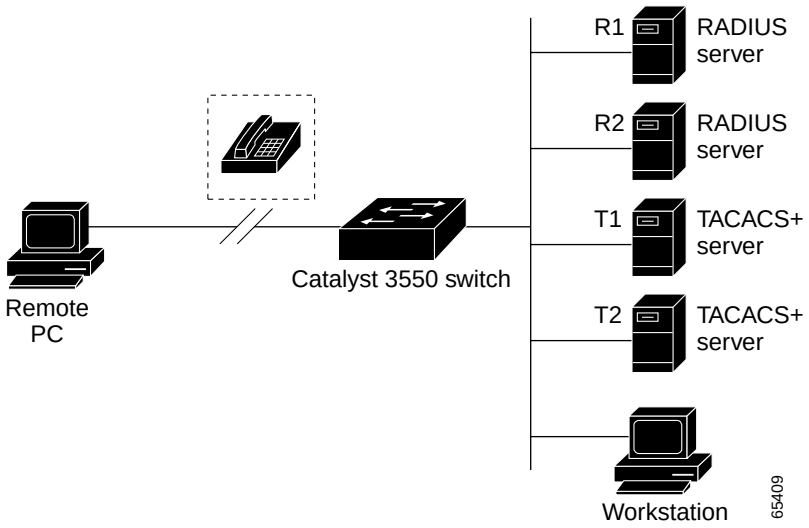

RADIUS Operation

When a user attempts to log in and authenticate to a switch that is access controlled by a RADIUS server, these events occur:

1. The user is prompted to enter a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of these responses from the RADIUS server:
   a. ACCEPT—The user is authenticated.
   b. REJECT—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
   c. CHALLENGE—A challenge requires additional data from the user.
   d. CHALLENGE PASSWORD—A response requests the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for privileged EXEC or network authorization. Users must first successfully complete RADIUS authentication before proceeding to RADIUS authorization, if it is enabled. The additional data included with the ACCEPT or REJECT packets includes these items:

- Telnet, rlogin, or privileged EXEC services
- Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring RADIUS

This section describes how to configure your switch to support RADIUS. At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used (such as TACACS+ or local username lookup), thus ensuring a backup system if the initial method fails. The

software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.

This section contains this configuration information:

- Default RADIUS Configuration, page 6-20
- Identifying the RADIUS Server Host, page 6-20 (required)
- Configuring RADIUS Login Authentication, page 6-23 (required)
- Defining AAA Server Groups, page 6-24 (optional)
- Configuring RADIUS Authorization for User Privileged Access and Network Services, page 6-26 (optional)
- Starting RADIUS Accounting, page 6-27 (optional)
- Configuring Settings for All RADIUS Servers, page 6-28 (optional)
- Configuring the Switch to Use Vendor-Specific RADIUS Attributes, page 6-28 (optional)
- Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, page 6-29 (optional)

Default RADIUS Configuration

RADIUS and AAA are disabled by default.

To prevent a lapse in security, you cannot configure RADIUS through a network management application. When enabled, RADIUS can authenticate users accessing the switch through the CLI.

Identifying the RADIUS Server Host

Switch-to-RADIUS-server communication involves several components:

- Host name or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their host name or IP address, host name and specific UDP port numbers, or their IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. This unique identifier enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address.

If two different host entries on the same RADIUS server are configured for the same service—for example, accounting—the second host entry configured acts as a fail-over backup to the first one. Using this example, if the first host entry fails to provide accounting services, the switch tries the second host entry configured on the same device for accounting services. (The RADIUS host entries are tried in the order that they are configured.)

A RADIUS server and the switch use a shared secret text string to encrypt passwords and exchange responses. To configure RADIUS to use the AAA security commands, you must specify the host running the RADIUS server daemon and a secret text (key) string that it shares with the switch.

The timeout, retransmission, and encryption key values can be configured globally for all RADIUS servers, on a per-server basis, or in some combination of global and per-server settings. To apply these settings globally to all RADIUS servers communicating with the switch, use the three unique global configuration commands: **radius-server timeout**, **radius-server retransmit**, and **radius-server key**. To apply these values on a specific RADIUS server, use the **radius-server host** global configuration command.

Note If you configure both global and per-server functions (timeout, retransmission, and key commands) on the switch, the per-server timer, retransmission, and key value commands override global timer, retransmission, and key value commands. For information on configuring these setting on all RADIUS servers, see the "Configuring Settings for All RADIUS Servers" section on page 6-28.

You can configure the switch to use AAA server groups to group existing server hosts for authentication. For more information, see the "Defining AAA Server Groups" section on page 6-24.

Beginning in privileged EXEC mode, follow these steps to configure per-server RADIUS server communication. This procedure is required.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} [**auth-port** *port-number*] [**acct-port** *port-number*] [**timeout** *seconds*] [**retransmit** *retries*] [**key** *string*]	Specify the IP address or host name of the remote RADIUS server host. • (Optional) For **auth-port** *port-number*, specify the UDP destination port for authentication requests. • (Optional) For **acct-port** *port-number*, specify the UDP destination port for accounting requests. • (Optional) For **timeout** *seconds*, specify the time interval that the switch waits for the RADIUS server to reply before retransmitting. The range is 1 to 1000. This setting overrides the **radius-server timeout** global configuration command setting. If no timeout is set with the **radius-server host** command, the setting of the **radius-server timeout** command is used. • (Optional) For **retransmit** *retries*, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the **radius-server host** command, the setting of the **radius-server retransmit** global configuration command is used. • (Optional) For **key** *string*, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the **radius-server host** command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key. To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the specified RADIUS server, use the **no radius-server host** *hostname* | *ip-address* global configuration command.

This example shows how to configure one RADIUS server to be used for authentication and another to be used for accounting:

```
Switch(config)# radius-server host 172.29.36.49 auth-port 1612 key rad1
Switch(config)# radius-server host 172.20.36.50 acct-port 1618 key rad2
```

This example shows how to configure *host1* as the RADIUS server and to use the default ports for both authentication and accounting:

```
Switch(config)# radius-server host host1
```

Note You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Configuring RADIUS Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication. This procedure is required.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa new-model**	Enable AAA.

	Command	Purpose
Step 3	**aaa authentication login** {**default** | *list-name*} *method1* [*method2...*]	Create a login authentication method list. • To create a default list that is used when a named list is *not* specified in the **login authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. • For *list-name*, specify a character string to name the list you are creating. • For *method1...*, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. Select one of these methods: • **line**—Use the line password for authentication. You must define a line password before you can use this authentication method. Use the **password** *password* line configuration command. • **local**—Use the local username database for authentication. You must enter username information in the database. Use the **username** *password* global configuration command. • **radius**—Use RADIUS authentication. You must configure the RADIUS server before you can use this authentication method. For more information, see the "Identifying the RADIUS Server Host" section on page 6-20.
Step 4	**line** [**console** | **tty** | **vty**] *line-number* [*ending-line-number*]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.
Step 5	**login authentication** {**default** | *list-name*}	Apply the authentication list to a line or set of lines. • If you specify **default**, use the default list created with the **aaa authentication login** command. • For *list-name*, specify the list created with the **aaa authentication login** command.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** {**default** | *list-name*} *method1* [*method2...*] global configuration command. To either disable RADIUS authentication for logins or to return to the default value, use the **no login authentication** {**default** | *list-name*} line configuration command.

Defining AAA Server Groups

You can configure the switch to use AAA server groups to group existing server hosts for authentication. You select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list, which lists the IP addresses of the selected server hosts.

Server groups also can include multiple host entries for the same server if each entry has a unique identifier (the combination of the IP address and UDP port number), allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. If you configure two different host entries on the same RADIUS server for the same service, (for example, accounting), the second configured host entry acts as a fail-over backup to the first one.

You use the **server** group server configuration command to associate a particular server with a defined group server. You can either identify the server by its IP address or identify multiple host instances or entries by using the optional **auth-port** and **acct-port** keywords.

Beginning in privileged EXEC mode, follow these steps to define the AAA server group and associate a particular RADIUS server with it:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} [**auth-port** *port-number*] [**acct-port** *port-number*] [**timeout** *seconds*] [**retransmit** *retries*] [**key** *string*]	Specify the IP address or host name of the remote RADIUS server host. • (Optional) For **auth-port** *port-number*, specify the UDP destination port for authentication requests. • (Optional) For **acct-port** *port-number*, specify the UDP destination port for accounting requests. • (Optional) For **timeout** *seconds*, specify the time interval that the switch waits for the RADIUS server to reply before retransmitting. The range is 1 to 1000. This setting overrides the **radius-server timeout** global configuration command setting. If no timeout is set with the **radius-server host** command, the setting of the **radius-server timeout** command is used. • (Optional) For **retransmit** *retries*, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the **radius-server host** command, the setting of the **radius-server retransmit** global configuration command is used. • (Optional) For **key** *string*, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the **radius-server host** command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key. To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.
Step 3	**aaa new-model**	Enable AAA.

	Command	Purpose
Step 4	**aaa group server radius** *group-name*	Define the AAA server-group with a group name. This command puts the switch in a server group configuration mode.
Step 5	**server** *ip-address*	Associate a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.
Step 9		Enable RADIUS login authentication. See the "Configuring RADIUS Login Authentication" section on page 6-23.

To remove the specified RADIUS server, use the **no radius-server host** *hostname | ip-address* global configuration command. To remove a server group from the configuration list, use the **no aaa group server radius** *group-name* global configuration command. To remove the IP address of a RADIUS server, use the **no server** *ip-address* server group configuration command.

In this example, the switch is configured to recognize two different RADIUS group servers (*group1* and *group2*). Group1 has two different host entries on the same RADIUS server configured for the same services. The second host entry acts as a fail-over backup to the first entry.

```
Switch(config)# radius-server host 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config)# radius-server host 172.10.0.1 auth-port 1645 acct-port 1646
Switch(config)# aaa new-model
Switch(config)# aaa group server radius group1
Switch(config-sg-radius)# server 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config-sg-radius)# exit
Switch(config)# aaa group server radius group2
Switch(config-sg-radius)# server 172.20.0.1 auth-port 2000 acct-port 2001
Switch(config-sg-radius)# exit
```

Configuring RADIUS Authorization for User Privileged Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **radius** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec radius local** command sets these authorization parameters:

- Use RADIUS for privileged EXEC access authorization if authentication was performed by using RADIUS.
- Use the local database if authentication was not performed by using RADIUS.

Note Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify RADIUS authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa authorization network radius**	Configure the switch for user RADIUS authorization for all network-related service requests.
Step 3	**aaa authorization exec radius**	Configure the switch for user RADIUS authorization to determine if the user has privileged EXEC access. The **exec** keyword might return user profile information (such as **autocommand** information).
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization** {**network** | **exec**} *method1* global configuration command.

Starting RADIUS Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable RADIUS accounting for each Cisco IOS privilege level and for network services:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa accounting network start-stop radius**	Enable RADIUS accounting for all network-related service requests.
Step 3	**aaa accounting exec start-stop radius**	Enable RADIUS accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting** {**network** | **exec**} {**start-stop**} *method1...* global configuration command.

Configuring Settings for All RADIUS Servers

Beginning in privileged EXEC mode, follow these steps to configure global communication settings between the switch and all RADIUS servers:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server key** *string*	Specify the shared secret text string used between the switch and all RADIUS servers. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 3	**radius-server retransmit** *retries*	Specify the number of times the switch sends each RADIUS request to the server before giving up. The default is 3; the range 1 to 1000.
Step 4	**radius-server timeout** *seconds*	Specify the number of seconds a switch waits for a reply to a RADIUS request before resending the request. The default is 5 seconds; the range is 1 to 1000.
Step 5	**radius-server deadtime** *minutes*	Specify the number of minutes a RADIUS server, which is not responding to authentication requests, to be skipped, thus avoiding the wait for the request to timeout before trying the next configured server. The default is 0; the range is 1 to 1440 minutes.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your settings.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting for the retransmit, timeout, and deadtime, use the **no** forms of these commands.

Configuring the Switch to Use Vendor-Specific RADIUS Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific information between the switch and the RADIUS server by using the vendor-specific attribute (attribute 26). Vendor-specific attributes (VSAs) allow vendors to support their own extended attributes not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option by using the format recommended in the specification. Cisco's vendor-ID is 9, and the supported option has vendor-type 1, which is named *cisco-avpair.* The value is a string with this format:

```
protocol : attribute sep value *
```

Protocol is a value of the Cisco protocol attribute for a particular type of authorization. *Attribute* and *value* are an appropriate attribute-value (AV) pair defined in the Cisco TACACS+ specification, and *sep* is = for mandatory attributes and * for optional attributes. This allows the full set of features available for TACACS+ authorization to also be used for RADIUS.

For example, the following AV pair activates Cisco's *multiple named ip address pools* feature during IP authorization (during PPP's IPCP address assignment):

```
cisco-avpair= "ip:addr-pool=first"
```

The following example shows how to provide a user logging in from a switch with immediate access to privileged EXEC commands:

```
cisco-avpair= "shell:priv-lvl=15"
```

Other vendors have their own unique vendor-IDs, options, and associated VSAs. For more information about vendor-IDs and VSAs, refer to RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)."

Beginning in privileged EXEC mode, follow these steps to configure the switch to recognize and use VSAs:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server vsa send** [**accounting** | **authentication**]	Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26. • (Optional) Use the **accounting** keyword to limit the set of recognized vendor-specific attributes to only accounting attributes. • (Optional) Use the **authentication** keyword to limit the set of recognized vendor-specific attributes to only authentication attributes. If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your settings.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

For a complete list of RADIUS attributes or more information about vendor-specific attribute 26, refer to the "RADIUS Attributes" appendix in the *Cisco IOS Security Configuration Guide for Release 12.1*.

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication

Although an IETF draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the switch and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. Cisco IOS software supports a subset of vendor-proprietary RADIUS attributes.

As mentioned earlier, to configure RADIUS (whether vendor-proprietary or IETF draft-compliant), you must specify the host running the RADIUS server daemon and the secret text string it shares with the switch. You specify the RADIUS host and secret text string by using the **radius-server** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to specify a vendor-proprietary RADIUS server host and a shared secret text string:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} **non-standard**	Specify the IP address or host name of the remote RADIUS server host and identify that it is using a vendor-proprietary implementation of RADIUS.
Step 3	**radius-server key** *string*	Specify the shared secret text string used between the switch and the vendor-proprietary RADIUS server. The switch and the RADIUS server use this text string to encrypt passwords and exchange responses. **Note** The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your settings.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete the vendor-proprietary RADIUS host, use the **no radius-server host** {*hostname* | *ip-address*} **non-standard** global configuration command. To disable the key, use the **no radius-server key** global configuration command.

This example shows how to specify a vendor-proprietary RADIUS host and to use a secret key of *rad124* between the switch and the server:

```
Switch(config)# radius-server host 172.20.30.15 nonstandard
Switch(config)# radius-server key rad124
```

Displaying the RADIUS Configuration

To display the RADIUS configuration, use the **show running-config** privileged EXEC command.

Configuring the Switch for Local Authentication and Authorization

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then handles authentication and authorization. No accounting is available in this configuration.

Beginning in privileged EXEC mode, follow these steps to configure the switch for local AAA:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa new-model**	Enable AAA.
Step 3	**aaa authentication login default local**	Set the login authentication to use the local username database. The **default** keyword applies the local user database authentication to all interfaces.
Step 4	**aaa authorization exec local**	Configure user AAA authorization to determine if the user is allowed to run an EXEC shell by checking the local database.
Step 5	**aaa authorization network local**	Configure user AAA authorization for all network-related service requests.
Step 6	**username** *name* [**privilege** *level*] {**password** *encryption-type password*}	Enter the local database, and establish a username-based authentication system. Repeat this command for each user. • For *name*, specify the user ID as one word. Spaces and quotation marks are not allowed. • (Optional) For *level*, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 0 gives user EXEC mode access. • For *encryption-type*, enter 0 to specify that an unencrypted password follows. Enter 7 to specify that a hidden password follows. • For *password*, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the **username** command.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show running-config**	Verify your entries.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable authorization, use the **no aaa authorization** {**network** | **exec**} *method1* global configuration command.

Managing the System Time and Date

You can manage the system time and date on your switch using automatic, such as the Network Time Protocol (NTP), or manual configuration methods.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- Understanding the System Clock, page 6-32
- Understanding Network Time Protocol, page 6-32
- Configuring NTP, page 6-34
- Configuring Time and Date Manually, page 6-41

Understanding the System Clock

The heart of the time service is the system clock. This clock runs from the moment the system starts up and keeps track of the current date and time.

The system clock can then be set from these sources:

- Network Time Protocol
- Manual configuration

The system clock can provide time to these services:

- User **show** commands
- Logging and debugging messages

The system clock keeps track of time internally based on Universal Time Coordinated (UTC), also known as Greenwich Mean Time (GMT). You can configure information about the local time zone and summer time (daylight saving time) so that the time is correctly displayed for the local time zone.

The system clock keeps track of whether the time is *authoritative* or not (that is, whether it has been set by a time source considered to be authoritative). If it is not authoritative, the time is available only for display purposes and is not redistributed. For configuration information, see the "Configuring Time and Date Manually" section on page 6-41.

Understanding Network Time Protocol

The NTP is designed to time-synchronize a network of devices. NTP runs over User Datagram Protocol (UDP), which runs over IP. NTP is documented in RFC 1305.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two devices to within a millisecond of one another.

NTP uses the concept of a *stratum* to describe how many NTP hops away a device is from an authoritative time source. A stratum 1 time server has a radio or atomic clock directly attached, a stratum 2 time server receives its time through NTP from a stratum 1 time server, and so on. A device

running NTP automatically chooses as its time source the device with the lowest stratum number with which it communicates through NTP. This strategy effectively builds a self-organizing tree of NTP speakers.

NTP avoids synchronizing to a device whose time might not be accurate by never synchronizing to a device that is not synchronized. NTP also compares the time reported by several devices and does not synchronize to a device whose time is significantly different than the others, even if its stratum is lower.

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP address of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, in that case, information flow is one-way only.

The time kept on a device is a critical resource; you should use the security features of NTP to avoid the accidental or malicious setting of an incorrect time. Two mechanisms are available: an access list-based restriction scheme and an encrypted authentication mechanism.

Cisco's implementation of NTP does not support stratum 1 service; it is not possible to connect to a radio or atomic clock. We recommend that the time service for your network be derived from the public NTP servers available on the IP Internet. Figure 6-3 show a typical network example using NTP.

If the network is isolated from the Internet, Cisco's implementation of NTP allows a device to act as though it is synchronized through NTP, when in fact it has determined the time by using other means. Other devices then synchronize to that device through NTP.

When multiple sources of time are available, NTP is always considered to be more authoritative. NTP time overrides the time set by any other method.

Several manufacturers include NTP software for their host systems, and a publicly available version for systems running UNIX and its various derivatives is also available. This software allows host systems to be time-synchronized as well.

Figure 6-3 Typical NTP Network Configuration

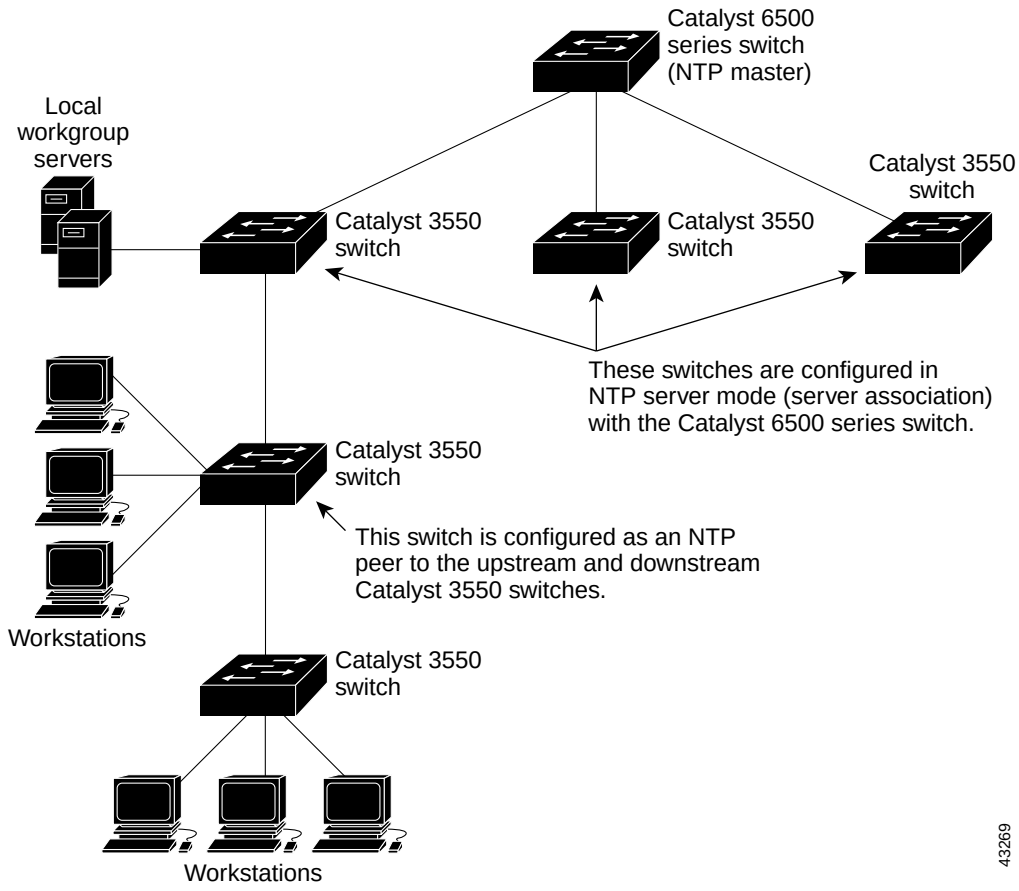

Configuring NTP

The Catalyst 3550 switches do not have a hardware-supported clock, and they cannot function as an NTP master clock to which peers synchronize themselves when an external NTP source is not available. These switches also have no hardware support for a calendar. As a result, the **ntp update-calendar** and the **ntp master** global configuration commands are not available.

This section contains this configuration information:

- Default NTP Configuration, page 6-35
- Configuring NTP Authentication, page 6-35
- Configuring NTP Associations, page 6-36
- Configuring NTP Broadcast Service, page 6-37
- Configuring NTP Access Restrictions, page 6-38
- Configuring the Source IP Address for NTP Packets, page 6-40
- Displaying the NTP Configuration, page 6-41

Default NTP Configuration

Table 6-2 shows the default NTP configuration.

Table 6-2 Default NTP Configuration

Feature	Default Setting
NTP authentication	Disabled. No authentication key is specified.
NTP peer or server associations	None configured.
NTP broadcast service	Disabled; no interface sends or receives NTP broadcast packets.
NTP access restrictions	No access control is specified.
NTP packet source IP address	The source address is determined by the outgoing interface.

NTP is enabled on all interfaces by default. All interfaces receive NTP packets.

Configuring NTP Authentication

This procedure must be coordinated with the administrator of the NTP server; the information you configure in this procedure must be matched by the servers used by the switch to synchronize its time to the NTP server.

Beginning in privileged EXEC mode, follow these steps to authenticate the associations (communications between devices running NTP that provide for accurate timekeeping) with other devices for security purposes:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ntp authenticate**	Enable the NTP authentication feature, which is disabled by default.
Step 3	**ntp authentication-key** *number* **md5** *value*	Define the authentication keys. By default, none are defined. • For *number*, specify a key number. The range is 1 to 4294967295. • **md5** specifies that message authentication support is provided by using the message digest algorithm 5 (MD5). • For *value*, enter an arbitrary string of up to eight characters for the key. The switch does not synchronize to a device unless both have one of these authentication keys, and the key number is specified by the **ntp trusted-key** *key-number* command.
Step 4	**ntp trusted-key** *key-number*	Specify one or more key numbers (defined in Step 3) that a peer NTP device must provide in its NTP packets for this switch to synchronize to it. By default, no trusted keys are defined. For *key-number*, specify the key defined in Step 3. This command provides protection against accidentally synchronizing the switch to a device that is not trusted.

	Command	Purpose
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable NTP authentication, use the **no ntp authenticate** global configuration command. To remove an authentication key, use the **no ntp authentication-key** *number* global configuration command. To disable authentication of the identity of a device, use the **no ntp trusted-key** *key-number* global configuration command.

This example shows how to configure the switch to synchronize only to devices providing authentication key 42 in the device's NTP packets:

```
Switch(config)# ntp authenticate
Switch(config)# ntp authentication-key 42 md5 aNiceKey
Switch(config)# ntp trusted-key 42
```

Configuring NTP Associations

An NTP association can be a peer association (this switch can either synchronize to the other device or allow the other device to synchronize to it), or it can be a server association (meaning that only this switch synchronizes to the other device, and not the other way around).

Beginning in privileged EXEC mode, follow these steps to form an NTP association with another device:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ntp peer** *ip-address* [**version** *number*] [**key** *keyid*] [**source** *interface*] [**prefer**]	Configure the switch system clock to synchronize a peer or to be synchronized by a peer (peer association).
	or	or
	ntp server *ip-address* [**version** *number*] [**key** *keyid*] [**source** *interface*] [**prefer**]	Configure the switch system clock to be synchronized by a time server (server association). No peer or server associations are defined by default. • For *ip-address* in a peer association, specify either the IP address of the peer providing, or being provided, the clock synchronization. For a server association, specify the IP address of the time server providing the clock synchronization. • (Optional) For *number*, specify the NTP version number. The range is 1 to 3. By default, version 3 is selected. • (Optional) For *keyid*, enter the authentication key defined with the **ntp authentication-key** global configuration command. • (Optional) For *interface*, specify the interface from which to pick the IP source address. By default, the source IP address is taken from the outgoing interface. • (Optional) Enter the **prefer** keyword to make this peer or server the preferred one that provides synchronization. This keyword reduces switching back and forth between peers and servers.

	Command	Purpose
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

You need to configure only one end of an association; the other device can automatically establish the association. If you are using the default NTP version (version 3) and NTP synchronization does not occur, try using NTP version 2. Many NTP servers on the Internet run version 2.

To remove a peer or server association, use the **no ntp peer** *ip-address* or the **no ntp server** *ip-address* global configuration command.

This example shows how to configure the switch to synchronize its system clock with the clock of the peer at IP address 172.16.22.44 using NTP version 2:

```
Switch(config)# ntp server 172.16.22.44 version 2
```

Configuring NTP Broadcast Service

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP addresses of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, the information flow is one-way only.

The switch can send or receive NTP broadcast packets on an interface-by-interface basis if there is an NTP broadcast server, such as a router, broadcasting time information on the network. The switch can send NTP broadcast packets to a peer so that the peer can synchronize to it. The switch can also receive NTP broadcast packets to synchronize its own clock. This section provides procedures for both sending and receiving NTP broadcast packets.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send NTP broadcast packets to peers so that they can synchronize their clock to the switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to send NTP broadcast packets.
Step 3	**ntp broadcast** [**version** *number*] [**key** *keyid*] [*destination-address*]	Enable the interface to send NTP broadcast packets to a peer. By default, this feature is disabled on all interfaces. • (Optional) For *number*, specify the NTP version number. The range is 1 to 3. If you do not specify a version, version 3 is used. • (Optional) For *keyid*, specify the authentication key to use when sending packets to the peer. • (Optional) For *destination-address*, specify the IP address of the peer that is synchronizing its clock to this switch.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.

	Command	Purpose
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.
Step 7		Configure the connected peers to receive NTP broadcast packets as described in the next procedure.

To disable the interface from sending NTP broadcast packets, use the **no ntp broadcast** interface configuration command.

This example shows how to configure an interface to send NTP version 2 packets:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ntp broadcast version 2
```

Beginning in privileged EXEC mode, follow these steps to configure the switch to receive NTP broadcast packets from connected peers:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to receive NTP broadcast packets.
Step 3	**ntp broadcast client**	Enable the interface to receive NTP broadcast packets. By default, no interfaces receive NTP broadcast packets.
Step 4	**exit**	Return to global configuration mode.
Step 5	**ntp broadcastdelay** *microseconds*	(Optional) Change the estimated round-trip delay between the switch and the NTP broadcast server. The default is 3000 microseconds; the range is 1 to 999999.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable an interface from receiving NTP broadcast packets, use the **no ntp broadcast client** interface configuration command. To change the estimated round-trip delay to the default, use the **no ntp broadcastdelay** global configuration command.

This example shows how to configure an interface to receive NTP broadcast packets:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ntp broadcast client
```

Configuring NTP Access Restrictions

You can control NTP access on two levels as described in these sections:

- Creating an Access Group and Assigning a Basic IP Access List, page 6-39
- Disabling NTP Services on a Specific Interface, page 6-40

Creating an Access Group and Assigning a Basic IP Access List

Beginning in privileged EXEC mode, follow these steps to control access to NTP services by using access lists:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ntp access-group** {**query-only** | **serve-only** | **serve** | **peer**} *access-list-number*	Create an access group, and apply a basic IP access list. The keywords have these meanings: • **query-only**—Allows only NTP control queries. • **serve-only**—Allows only time requests. • **serve**—Allows time requests and NTP control queries, but does not allow the switch to synchronize to the remote device. • **peer**—Allows time requests and NTP control queries and allows the switch to synchronize to the remote device. For *access-list-number*, enter a standard IP access list number from 1 to 99.
Step 3	**access-list** *access-list-number* **permit** *source* [*source-wildcard*]	Create the access list. • For *access-list-number*, enter the number specified in Step 2. • Enter the **permit** keyword to permit access if the conditions are matched. • For *source*, enter the IP address of the device that is permitted access to the switch. • (Optional) For *source-wildcard*, enter the wildcard bits to be applied to the source. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The access group keywords are scanned in this order, from least restrictive to most restrictive:

1. **peer**—Allows time requests and NTP control queries and allows the switch to synchronize itself to a device whose address passes the access list criteria.
2. **serve**—Allows time requests and NTP control queries, but does not allow the switch to synchronize itself to a device whose address passes the access list criteria.
3. **serve-only**—Allows only time requests from a device whose address passes the access list criteria.
4. **query-only**—Allows only NTP control queries from a device whose address passes the access list criteria.

If the source IP address matches the access lists for more than one access type, the first type is granted. If no access groups are specified, all access types are granted to all devices. If any access groups are specified, only the specified access types are granted.

To remove access control to the switch NTP services, use the **no ntp access-group** {**query-only** | **serve-only** | **serve** | **peer**} global configuration command.

This example shows how to configure the switch to allow itself to synchronize to a peer from access list 99. However, the switch restricts access to allow only time requests from access list 42:

```
Switch# configure terminal
Switch(config)# ntp access-group peer 99
Switch(config)# ntp access-group serve-only 42
Switch(config)# access-list 99 permit 172.20.130.5
Switch(config)# access list 42 permit 172.20.130.6
```

Disabling NTP Services on a Specific Interface

NTP services are enabled on all interfaces by default.

Beginning in privileged EXEC mode, follow these steps to disable NTP packets from being received on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to disable.
Step 3	**ntp disable**	Disable NTP packets from being received on the interface. By default, all interfaces receive NTP packets.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To re-enable receipt of NTP packets on an interface, use the **no ntp disable** interface configuration command.

Configuring the Source IP Address for NTP Packets

When the switch sends an NTP packet, the source IP address is normally set to the address of the interface through which the NTP packet is sent. Use the **ntp source** global configuration command when you want to use a particular source IP address for all NTP packets. The address is taken from the specified interface. This command is useful if the address on an interface cannot be used as the destination for reply packets.

Beginning in privileged EXEC mode, follow these steps to configure a specific interface from which the IP source address is to be taken:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ntp source** *type number*	Specify the interface type and number from which the IP source address is taken. By default, the source address is determined by the outgoing interface.

	Command	Purpose
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The specified interface is used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the **source** keyword in the **ntp peer** or **ntp server** global configuration command as described in the "Configuring NTP Associations" section on page 6-36.

Displaying the NTP Configuration

You can use two privileged EXEC commands to display NTP information:

- **show ntp associations** [**detail**]
- **show ntp status**

For detailed information about the fields in these displays, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

Configuring Time and Date Manually

If no other source of time is available, you can manually configure the current time and date after the system is restarted. The time remains accurate until the next system restart. We recommend that you use manual configuration only as a last resort. If you have an outside source to which the switch can synchronize, you do not need to manually set the system clock.

This section contains this configuration information:

- Setting the System Clock, page 6-42
- Displaying the Time and Date Configuration, page 6-42
- Configuring the Time Zone, page 6-43
- Configuring Summer Time (Daylight Saving Time), page 6-44

Setting the System Clock

If you have an outside source on the network that provides time services, such as an NTP server, you do not need to manually set the system clock.

Beginning in privileged EXEC mode, follow these steps to set the system clock:

	Command	Purpose
Step 1	**clock set** *hh*:*mm*:*ss day month year* or **clock set** *hh*:*mm*:*ss month day year*	Manually set the system clock using one of these formats. • For *hh*:*mm*:*ss*, specify the current time in hours (24-hour format), minutes, and seconds. The time specified is relative to the configured time zone. • For *day*, specify the day by date in the month. • For *month*, specify the month by name. • For *year*, specify the year (no abbreviation).
Step 2	**show running-config**	Verify your entries.
Step 3	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

This example shows how to manually set the system clock to 1:32 p.m. on July 23, 2001:

```
Switch# clock set 13:32:00 23 July 2001
```

Displaying the Time and Date Configuration

To display the time and date configuration, use the **show clock** [**detail**] privileged EXEC command.

The system clock keeps an *authoritative* flag that shows whether the time is authoritative (believed to be accurate). If the system clock has been set by a timing source such as NTP, the flag is set. If the time is not authoritative, it is used only for display purposes. Until the clock is authoritative and the *authoritative* flag is set, the flag prevents peers from synchronizing to the clock when the peers' time is invalid.

The symbol that precedes the **show clock** display has this meaning:

- *—Time is not authoritative.
- (blank)—Time is authoritative.
- .—Time is authoritative, but NTP is not synchronized.

Configuring the Time Zone

Beginning in privileged EXEC mode, follow these steps to manually configure the time zone:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**clock timezone** *zone hours-offset* [*minutes-offset*]	Set the time zone. The switch keeps internal time in universal time coordinated (UTC), so this command is used only for display purposes and when the time is manually set. • For *zone*, enter the name of the time zone to be displayed when standard time is in effect. The default is UTC. • For *hours-offset*, enter the hours offset from UTC. • (Optional) For *minutes-offset*, enter the minutes offset from UTC.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The *minutes-offset* variable in the **clock timezone** global configuration command is available for those cases where a local time zone is a percentage of an hour different from UTC. For example, the time zone for some sections of Atlantic Canada (AST) is UTC-3.5, where the 3 means 3 hours and .5 means 50 percent. In this case, the necessary command is **clock timezone AST -3 30**.

To set the time to UTC, use the **no clock timezone** global configuration command.

Configuring Summer Time (Daylight Saving Time)

Beginning in privileged EXEC mode, follow these steps to configure summer time (daylight saving time) in areas where it starts and ends on a particular day of the week each year:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**clock summer-time** *zone* **recurring** [*week day month hh*:*mm week day month hh*:*mm* [*offset*]]	Configure summer time to start and end on the specified days every year. Summer time is disabled by default. If you specify **clock summer-time** *zone* **recurring** without parameters, the summer time rules default to the United States rules. • For *zone*, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. • (Optional) For *week*, specify the week of the month (1 to 5 or **last**). • (Optional) For *day*, specify the day of the week (Sunday, Monday...). • (Optional) For *month*, specify the month (January, February...). • (Optional) For *hh*:*mm*, specify the time (24-hour format) in hours and minutes. • (Optional) For *offset*, specify the number of minutes to add during summer time. The default is 60.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

This example shows how to specify that summer time starts on the first Sunday in April at 02:00 and ends on the last Sunday in October at 02:00:

```
Switch(config)# clock summer-time PDT recurring 1 Sunday April 2:00 last Sunday October
2:00
```

Beginning in privileged EXEC mode, follow these steps if summer time in your area does not follow a recurring pattern (configure the exact date and time of the next summer time events):

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**clock summer-time** *zone* **date** [*month date year hh*:*mm month date year hh*:*mm* [*offset*]] or **clock summer-time** *zone* **date** [*date month year hh*:*mm date month year hh*:*mm* [*offset*]]	Configure summer time to start on the first date and end on the second date. Summer time is disabled by default. • For *zone*, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. • (Optional) For *week*, specify the week of the month (1 to 5 or **last**). • (Optional) For *day*, specify the day of the week (Sunday, Monday...). • (Optional) For *month*, specify the month (January, February...). • (Optional) For *hh*:*mm*, specify the time (24-hour format) in hours and minutes. • (Optional) For *offset*, specify the number of minutes to add during summer time. The default is 60.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

To disable summer time, use the **no clock summer-time** global configuration command.

This example shows how to set summer time to start on October 12, 2000, at 02:00, and end on April 26, 2001, at 02:00:

```
Switch(config)# clock summer-time pdt date 12 October 2000 2:00 26 April 2001 2:00
```

Configuring a System Name and Prompt

You configure the system name on the switch to identify it. By default, the system name and prompt are *Switch*.

If you have not configured a system prompt, the first 20 characters of the system name are used as the system prompt. A greater-than symbol [>] is appended. The prompt is updated whenever the system name changes, unless you manually configure the prompt by using the **prompt** global configuration command.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference* and the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section contains this configuration information:

- Default System Name and Prompt Configuration, page 6-46
- Configuring a System Name, page 6-46
- Configuring a System Prompt, page 6-47
- Understanding DNS, page 6-47

Default System Name and Prompt Configuration

The default switch system name and prompt is *Switch*.

Configuring a System Name

Beginning in privileged EXEC mode, follow these steps to manually configure a system name:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**hostname** *name*	Manually configure a system name. The default setting is *switch*. The name must follow the rules for ARPANET host names. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphens. Names can be up to 63 characters.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

When you set the system name, it is also used as the system prompt. You can override the prompt setting by using the **prompt** global configuration command.

To return to the default hostname, use the **no hostname** global configuration command.

Configuring a System Prompt

Beginning in privileged EXEC mode, follow these steps to manually configure a system prompt:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**prompt** *string*	Configure the command-line prompt to override the setting from the **hostname** command. The default prompt is either *switch* or the name defined with the **hostname** global configuration command, followed by an angle bracket (>) for user EXEC mode or a pound sign (#) for privileged EXEC mode. The prompt can consist of all printing characters and escape sequences.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default prompt, use the **no prompt** [*string*] global configuration command.

Understanding DNS

The DNS protocol controls the Domain Name System (DNS), a distributed database with which you can map host names to IP addresses. When you configure DNS on your switch, you can substitute the host name for the IP address with all IP commands, such as **ping**, **telnet**, **connect**, and related Telnet support operations.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that IP identifies by a *com* domain name, so its domain name is *cisco.com*. A specific device in this domain, for example, the File Transfer Protocol (FTP) system is identified as *ftp.cisco.com*.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, specify the name server that is present on your network, and enable the DNS.

This section contains this configuration information:

- Default DNS Configuration, page 6-48
- Setting Up DNS, page 6-48
- Displaying the DNS Configuration, page 6-49

Default DNS Configuration

Table 6-3 shows the default DNS configuration.

Table 6-3 Default DNS Configuration

Feature	Default Setting
DNS enable state	Enabled.
DNS default domain name	None configured.
DNS servers	No name server addresses are configured.

Setting Up DNS

Beginning in privileged EXEC mode, follow these steps to set up your switch to use the DNS:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip domain-name** *name*	Define a default domain name that the software uses to complete unqualified host names (names without a dotted-decimal domain name). Do not include the initial period that separates an unqualified name from the domain name. At boot time, no domain name is configured; however, if the switch configuration comes from a BOOTP or Dynamic Host Configuration Protocol (DHCP) server, then the default domain name might be set by the BOOTP or DHCP server (if the servers were configured with this information).
Step 3	**ip name-server** *server-address1* [*server-address2 ... server-address6*]	Specify the address of one or more name servers to use for name and address resolution. You can specify up to six name servers. Separate each server address with a space. The first server specified is the primary server. The switch sends DNS queries to the primary server first. If that query fails, the backup servers are queried.
Step 4	**ip domain-lookup**	(Optional) Enable DNS-based host name-to-address translation on your switch. This feature is enabled by default. If your network devices require connectivity with devices in networks for which you do not control name assignment, you can dynamically assign device names that uniquely identify your devices by using the global Internet naming scheme (DNS).
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

If you use the switch IP address as its hostname, the IP address is used and no DNS query occurs. If you configure a hostname that contains no periods (.), a period followed by the default domain name is appended to the hostname before the DNS query is made to map the name to an IP address. The default

domain name is the value set by the **ip domain-name** global configuration command. If there is a period (.) in the hostname, the IOS software looks up the IP address without appending any default domain name to the hostname.

To remove a domain name, use the **no ip domain-name** *name* global configuration command. To remove a name server address, use the **no ip name-server** *server-address* global configuration command. To disable DNS on the switch, use the **no ip domain-lookup** global configuration command.

Displaying the DNS Configuration

To display the DNS configuration information, use the **show running-config** privileged EXEC command.

Creating a Banner

You can configure a message-of-the-day (MOTD) and a login banner. The MOTD banner displays on all connected terminals at login and is useful for sending messages that affect all network users (such as impending system shutdowns).

The login banner also displays on all connected terminals. It is displayed after the MOTD banner and before the login prompts.

Note For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- Default Banner Configuration, page 6-49
- Configuring a Message-of-the-Day Login Banner, page 6-50
- Configuring a Login Banner, page 6-51

Default Banner Configuration

The MOTD and login banners are not configured.

Configuring a Message-of-the-Day Login Banner

You can create a single or multiline message banner that appears on the screen when someone logs in to the switch.

Beginning in privileged EXEC mode, follow these steps to configure a MOTD login banner:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**banner motd** *c message c*	Specify the message of the day. For *c*, enter the delimiting character of your choice, for example, a pound sign (#), and press the **Return** key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For *message*, enter a banner message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete the MOTD banner, use the **no banner motd** global configuration command.

This example shows how to configure a MOTD banner for the switch using the pound sign (#) symbol as the beginning and ending delimiter:

```
Switch(config)# banner motd #
This is a secure site. Only authorized users are allowed.
For access, contact technical support.
#
Switch(config)#
```

This example shows the banner displayed from the previous configuration:

```
Unix> telnet 172.2.5.4
Trying 172.2.5.4...
Connected to 172.2.5.4.
Escape character is '^]'.

This is a secure site. Only authorized users are allowed.
For access, contact technical support.

User Access Verification

Password:
```

Configuring a Login Banner

You can configure a login banner to be displayed on all connected terminals. This banner is displayed after the MOTD banner and before the login prompt.

Beginning in privileged EXEC mode, follow these steps to configure a login banner:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**banner login** *c message c*	Specify the login message. For *c*, enter the delimiting character of your choice, for example, a pound sign (#), and press the **Return** key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For *message*, enter a login message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete the login banner, use the **no banner login** global configuration command.

This example shows how to configure a login banner for the switch using the dollar sign ($) symbol as the beginning and ending delimiter:

```
Switch(config)# banner login $
Access for authorized users only. Please enter your username and password.
$
Switch(config)#
```

Managing the MAC Address Table

The MAC address table contains address information that the switch uses to forward traffic between ports. All MAC addresses in the address table are associated with one or more ports. The address table includes these types of addresses:

- Dynamic address: a source MAC address that the switch learns and then ages when it is not in use.
- Static address: a manually entered unicast or multicast address that does not age and that is not lost when the switch resets.

The address table lists the destination MAC address, the associated VLAN ID, and port number associated with the address.

Note For complete syntax and usage information for the commands used in this section, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This section contains this configuration information:

- Building the Address Table, page 6-52
- MAC Addresses and VLANs, page 6-52
- Default MAC Address Table Configuration, page 6-53
- Changing the Address Aging Time, page 6-53
- Removing Dynamic Address Entries, page 6-54
- Configuring MAC Address Notification Traps, page 6-54
- Adding and Removing Static Address Entries, page 6-56
- Displaying Address Table Entries, page 6-57

Building the Address Table

With multiple MAC addresses supported on all ports, you can connect any port on the switch to individual workstations, repeaters, switches, routers, or other network devices. The switch provides dynamic addressing by learning the source address of packets it receives on each port and adding the address and its associated port number to the address table. As stations are added or removed from the network, the switch updates the address table, adding new dynamic addresses and aging out those that are currently not in use.

The aging interval is configured on a per-switch basis. However, the switch maintains an address table for each VLAN, and STP can accelerate the aging interval on a per-VLAN basis.

The switch sends packets between any combination of ports, based on the destination address of the received packet. Using the MAC address table, the switch forwards the packet only to the port or ports associated with the destination address. If the destination address is on the port that sent the packet, the packet is filtered and not forwarded. The switch always uses the store-and-forward method: complete packets are stored and checked for errors before transmission.

MAC Addresses and VLANs

All addresses are associated with a VLAN. An address can exist in more than one VLAN and have different destinations in each. Multicast addresses, for example, could be forwarded to port 1 in VLAN 1 and ports 9, 10, and 11 in VLAN 5.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN. Addresses that are statically entered in one VLAN must be configured as static addresses in all other VLANs or remain unlearned in the other VLANs.

Default MAC Address Table Configuration

Table 6-4 shows the default MAC address table configuration.

Table 6-4 Default MAC Address Table Configuration

Feature	Default Setting
Aging time	300 seconds
Dynamic addresses	Automatically learned
Static addresses	None configured

Changing the Address Aging Time

Dynamic addresses are source MAC addresses that the switch learns and then ages when they are not in use. You can change the aging time setting for all VLANs or for a specified VLAN.

Setting too short an aging time can cause addresses to be prematurely removed from the table. Then when the switch receives a packet for an unknown destination, it floods the packet to all ports in the same VLAN as the receiving port. This unnecessary flooding can impact performance. Setting too long an aging time can cause the address table to be filled with unused addresses, which prevents new addresses from being learned. Flooding results, which can impact switch performance.

Beginning in privileged EXEC mode, follow these steps to configure the dynamic address table aging time:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mac-address-table aging-time** [**0** | *10-1000000*] [**vlan** *vlan-id*]	Set the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. The range is 10 to 1000000 seconds. The default is 300. You can also enter 0, which disables aging. Static address entries are never aged or removed from the table. For *vlan-id*, valid IDs are 1 to 1005.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mac-address-table aging-time**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no mac-address-table aging-time** global configuration command.

Removing Dynamic Address Entries

To remove all dynamic entries, use the **clear mac-address-table dynamic** command in privileged EXEC mode. You can also remove a specific MAC address (**clear mac-address-table dynamic address** *mac-address*), remove all addresses on the specified physical port or port channel (**clear mac-address-table dynamic interface** *interface-id*), or remove all addresses on a specified VLAN (**clear mac-address-table dynamic vlan** *vlan-id*).

To verify that dynamic entries have been removed, use the **show mac-address-table dynamic** privileged EXEC command.

Configuring MAC Address Notification Traps

MAC address notification enables you to track users on a network by storing the MAC address activity on the switch. Whenever the switch learns or removes a MAC address, an SNMP notification can be generated and sent to the NMS. If you have many users coming and going from the network, you can set a trap interval time to bundle the notification traps and reduce network traffic. The MAC notification history table stores the MAC address activity for each hardware port for which the trap is enabled. MAC address notifications are generated for dynamic and secure MAC addresses; events are not generated for self addresses, multicast addresses, or other static addresses.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send MAC address notification traps to an NMS host:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**snmp-server host** *host-addr* {**traps** | **informs**} {**version** {**1** | **2c**}} *community-string notification-type*	Specify the recipient of the trap message. • For *host-addr*, specify the name or address of the NMS. • Specify **traps** (the default) to send SNMP traps to the host. Specify **informs** to send SNMP informs to the host. • Specify the SNMP version to support. Version 1, the default, is not available with informs. **Note** Though visible in the command-line help string, the **version 3** keyword (SNMPv3) is not supported. • For *community-string*, specify the string to send with the notification operation. Though you can set this string using the **snmp-server host** command, we recommend that you define this string by using the **snmp-server community** command before using the **snmp-server host** command. • For *notification-type*, use the **mac-notification** keyword.

	Command	Purpose
Step 3	**snmp-server enable traps mac-notification**	Enable the switch to send MAC address traps to the NMS.
Step 4	**mac-address-table notification**	Enable the MAC address notification feature.
Step 5	**mac-address-table notification** [**interval** *value*] | [**history-size** *value*]	Enter the trap interval time and the history table size. • (Optional) For **interval** *value*, specify the notification trap interval in seconds between each set of traps that are generated to the NMS. The range is 0 to 2147483647 seconds; the default is 1 second. • (Optional) For **history-size** *value*, specify the maximum number of entries in the MAC notification history table. The range is 0 to 500; the default is 1.
Step 6	**interface** *interface-id*	Enter interface configuration mode, and specify the interface on which to enable the snmp MAC address notification trap.
Step 7	**snmp trap mac-notification** {**added** | **removed**}	Enable the MAC address notification trap. • Enable the MAC notification trap whenever a MAC address is **added** on this interface. • Enable the MAC notification trap whenever a MAC address is **removed** from this interface.
Step 8	**end**	Return to privileged EXEC mode.
Step 9	**show mac-address-table notification interface** **show running-config**	Verify your entries.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the switch from sending MAC address notification traps, use the **no snmp-server enable traps mac-notification** global configuration command. To disable the MAC address notification traps on a specific interface, use the **no snmp trap mac-notification** {**added** | **removed**} interface configuration command. To disable the MAC address notification feature, use the **no mac-address-table notification** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address notification traps to the NMS, enable the MAC address notification feature, set the interval time to 60 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on Fast Ethernet interface 0/4.

```
Switch(config)# snmp-server host 172.20.10.10 traps private
Switch(config)# snmp-server enable traps mac-notification
Switch(config)# mac-address-table notification
Switch(config)# mac-address-table notification interval 60
Switch(config)# mac-address-table notification history-size 100
Switch(config)# interface fastethernet0/4
Switch(config-if)# snmp trap mac-notification added
```

You can verify the previous commands by entering the **show mac-address-table notification interface** and the **show mac-address-table notification** privileged EXEC commands.

Adding and Removing Static Address Entries

A static address has these characteristics:

- It is manually entered in the address table and must be manually removed.
- It can be a unicast or multicast address.
- It does not age and is retained when the switch restarts.

You can add and remove static addresses and define the forwarding behavior for them. The forwarding behavior determines how a port that receives a packet forwards it to another port for transmission. Because all ports are associated with at least one VLAN, the switch acquires the VLAN ID for the address from the ports that you specify. You can specify a different list of destination ports for each source port.

A static address in one VLAN must be a static address in other VLANs. A packet with a static address that arrives on a VLAN where it has not been statically entered is flooded to all ports and not learned.

You add a static address to the address table by specifying the destination MAC address (unicast or multicast) and the VLAN from which it is received. Packets received with this destination address are forwarded to the interface specified with the *interface-id* option.

Beginning in privileged EXEC mode, follow these steps to add a static address:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mac-address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id*	Add a static address to the MAC address table. • For *mac-addr*, specify the destination MAC address (unicast or multicast) to add to the address table. Packets with this destination address received in the specified VLAN are forwarded to the specified interface. • For *vlan-id*, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 1005; do not enter leading zeros. • For *interface-id...*, specify the interface to which the received packet is forwarded. Valid interfaces include physical ports.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mac-address-table static**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove static entries from the address table, use the **no mac-address-table static** *mac-addr* **vlan** *vlan-id* **interface** *interface-id* global configuration command.

This example shows how to add the static address c2f3.220a.12f4 to the MAC address table. When a packet is received in VLAN 4 with this MAC address as its destination address, the packets is forwarded to the specified interface:

```
Switch(config)# mac-address-table static c2f3.220a.12f4 vlan 4 interface
gigabitethernet0/1
```

Displaying Address Table Entries

You can display the MAC address table by using one or more of the privileged EXEC commands described in Table 6-5:

Table 6-5 Commands for Displaying the MAC Address Table

Command	Description
show mac-address-table address	Displays MAC address table information for the specified MAC address.
show mac-address-table aging-time	Displays the aging time in all VLANs or the specified VLAN.
show mac-address-table count	Displays the number of addresses present in all VLANs or the specified VLAN.
show mac-address-table dynamic	Displays dynamic MAC address table entries only.
show mac-address-table interface	Displays the MAC address table information for the specified interface.
show mac-address-table multicast	Displays the Layer 2 multicast entries for all VLANs or the specified VLAN.
show mac-address-table static	Displays static MAC address table entries only.
show mac-address-table vlan	Displays the MAC address table information for the specified VLAN.

Optimizing System Resources for User-Selected Features

By using Switch Database Management (SDM) templates, you can configure memory resources in the switch to optimize support for specific features, depending on how the switch is used in your network. You can select one of four templates to specify how system resources are allocated. You can then approximate the maximum number of unicast MAC addresses, Internet Group Management Protocol (IGMP) groups, quality of service (QoS) access control entries (ACEs), security ACEs, unicast routes, multicast routes, subnet VLANs (routed interfaces), and Layer 2 VLANs that can be configured on the switch.

The four templates prioritize system memory to optimize support for these types of features:

- QoS and security ACEs—The access template might typically be used in an access switch at the network edge where the route table sizes might not be substantial. Filtering and QoS might be more important because an access switch is the entry to the whole network.
- Routing—The routing template maximizes system resources for unicast routing, typically required for a router or aggregator in the center of a network.
- VLANs—The VLAN template disables routing and supports the maximum number of unicast MAC addresses. It would typically be selected for a Catalyst 3550 used as a Layer 2 switch.
- Default—The default template gives balance to all functionalities (QoS, ACLs, unicast routing, multicast routing, VLANs and MAC addresses).

Table 6-6 lists the approximate number of each resource supported in each of the four templates for Catalyst 3550 Gigabit Ethernet switches. Table 6-7 compares the four templates for a Catalyst 3550 switch with primarily Fast Ethernet ports.

The first six rows in the tables (unicast MAC addresses through multicast routes) represent approximate hardware boundaries set when a template is selected. If a section of a hardware resource is full, all processing overflow is sent to the CPU, seriously impacting switch performance.

The last two rows, the total number of routed ports and SVIs and the number of Layer 2 VLANs, are guidelines used to calculate hardware resource consumption related to the other resource parameters.

The number of subnet VLANs (routed ports and SVIs) are not limited by software and can be set to a number higher than indicated in the tables. If the number of subnet VLANs configured is lower or equal to the number in the tables, the number of entries in each category (unicast addresses, IGMP groups, and so on) for each template will be as shown. As the number of subnet VLANs increases, CPU utilization typically increases. If the number of subnet VLANs increases beyond the number shown in the tables, the number of supported entries in each category could decrease depending on features that are enabled. For example, if PIM-DVMRP is enabled with more than 16 subnet VLANs, the number of entries for multicast routes will be in the range of 1K-5K entries for the access template.

Table 6-6 Approximate Resources Allowed in Each Template for Gigabit Ethernet Switches

Resource	Default Template	Access Template	Routing Template	VLAN Template
Unicast MAC addresses	6 K	2 K	6 K	12 K
IGMP groups (managed by Layer 2 multicast features such as MVR or IGMP snooping)	6 K	8 K	6 K	6 K
QoS classification ACEs	2 K	2 K	1 K	2 K
Security ACEs	2 K	4 K	1 K	2 K
Unicast routes	12 K	4 K	24 K	0
Multicast routes	6 K	8 K	6 K	0
Subnet VLANs (routed ports and SVIs)	16	16	16	16
Layer 2 VLANs	1 K	1 K	1 K	1 K

Table 6-7 Approximate Resources Allowed in Each Template for Fast Ethernet Switches

Resource	Default Template	Access Template	Routing Template	VLAN Template
Unicast MAC addresses	5 K	1 K	5 K	8 K
IGMP groups (managed by Layer 2 multicast features such as MVR and IGMP snooping)	1 K	2 K	1 K	1 K
QoS classification ACEs	1 K	2 K	512	1 K
Security ACEs	1 K	2 K	512	1 K
Unicast routes	8 K	2 K	16 K	0
Multicast routes	1 K	2 K	1 K	0
Subnet VLANs (routed ports and SVIs)	8	8	8	8
Layer 2 VLANs	1 K	1 K	1 K	1 K

Using the Templates

Follow these guidelines when using the SDM templates:

- The maximum number of resources allowed in each template is an approximation and depends upon the actual number of other features configured. For example, in the default template for the Catalyst 3550-12T, if your switch has more than 16 routed interfaces configured, the number of multicast or unicast routes that can be accommodated by hardware might be fewer than shown.
- Using the **sdm prefer vlan** global configuration command disables routing capability in the switch. Any routing configurations will be rejected after the reload, and those options might be lost. When the switch reloads, the message "Running Layer2 Switching Only Image" appears, even though you are really running a Layer 3 image with the routing functionality disabled. Use the **sdm prefer vlan** global configuration command only on switches intended for Layer 2 switching with no routing.
- Do not use the routing template if you do not have the enhanced multilayer software image installed on your switch. Entering the **sdm prefer routing** global configuration command on a switch that does not have the enhanced multilayer software image would not enable routing, but it would prevent other features from using the approximately 17 K of memory allocated to unicast and multicast routing in the routing template.

This procedure shows how to change the SDM template from the default. The switch must reload before the configuration takes effect. If you use the **show sdm prefer** privileged EXEC command before the switch reloads, the previous configuration (in this case, the default) is displayed.

Beginning in privileged EXEC mode, follow these steps to use the SDM template to maximize feature usage:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**sdm prefer** {**access** | **routing** | **vlan**}	Specify the SDM template to be used on the switch: The keywords have these meanings: • **access**—Maximizes the use of QoS classification ACEs and security ACEs on the switch. • **routing**—Maximizes routing on the switch. • **vlan**—Maximizes VLAN configuration on the switch with no routing allowed. The default template (if none of these is configured) balances the use of unicast MAC addresses, IGMP groups, QoS ACEs, security ACEs, unicast and multicast routes, routed interfaces, and Layer 2 VLANs.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**copy running-config startup-config**	Save your entries in the configuration file.
Step 5	**reload**	Reload the operating system.

After the system reboots, you can use the **show sdm prefer** privileged EXEC command to verify the new template configuration. If you use the **show sdm prefer** command before the **reload** privileged EXEC command, the previous template is displayed instead of the new one.

To return to the default template, use the **no sdm prefer** {**access** | **routing** | **vlan**} global configuration command.

This example shows how to configure a switch with the routing template and verify the configuration:

```
Switch(config)# sdm prefer routing
Switch(config)# end
Switch# copy running-config startup-config
Switch# reload
Proceed with reload? [confirm]


Switch# show sdm prefer
 The current template is routing template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 16 routed interfaces and 1K VLANs.

 number of unicast mac addresses:  6K
 number of igmp groups:            6K
 number of qos aces:               1K
 number of security aces:          1K
 number of unicast routes:         24K
 number of multicast routes:       6K
```

CHAPTER 7

Configuring 802.1X Port-Based Authentication

This chapter describes how to configure IEEE 802.1X port-based authentication to prevent unauthorized devices (clients) from gaining access to the network. As LANs extend to hotels, airports, and corporate lobbies, insecure environments could be created.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding 802.1X Port-Based Authentication, page 7-1
- Configuring 802.1X Authentication, page 7-5
- Displaying 802.1X Statistics and Status, page 7-14

Understanding 802.1X Port-Based Authentication

The IEEE 802.1X standard defines a client-server-based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port before making available any services offered by the switch or the LAN.

Until the client is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.

This section includes this conceptual information:

- Device Roles, page 7-2
- Authentication Initiation and Message Exchange, page 7-3
- Ports in Authorized and Unauthorized States, page 7-4
- Supported Topologies, page 7-4

Device Roles

With 802.1X port-based authentication, the devices in the network have specific roles as shown in Figure 7-1.

Figure 7-1 802.1X Device Roles

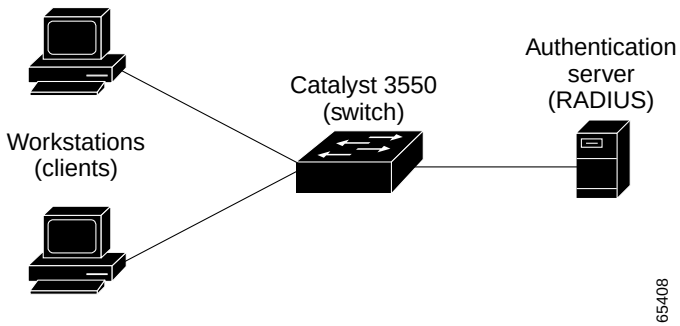

- *Client*—the device (workstation) that requests access to the LAN and switch services and responds to requests from the switch.The workstation must be running 802.1X-compliant client software such as that offered in the Microsoft Windows XP operating system. (The client is the *supplicant* in the IEEE 802.1X specification.)

  Note To resolve Windows XP network connectivity and 802.1X authentication issues, read the Microsoft Knowledge Base article at this URL: http://support.microsoft.com/support/kb/articles/Q303/5/97.ASP

- *Authentication server*—performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server; it is available in Cisco Secure Access Control Server version 3.0. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
- *Switch* (edge switch or wireless access point)—controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the Extensible Authentication Protocol (EAP) frames and interacting with the authentication server.

  When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is re-encapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

  The devices that can act as intermediaries include the Catalyst 3550 multilayer switch, the Catalyst 2950 switch, or a wireless access point. These devices must be running software that supports the RADIUS client and 802.1X.

Authentication Initiation and Message Exchange

The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link state transitions from down to up. It then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during bootup, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.

Note If 802.1X is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client transmits frames as if the port is in the authorized state. A port in the authorized state effectively means that the client has been successfully authenticated. For more information, see the "Ports in Authorized and Unauthorized States" section on page 7-4.

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized. For more information, see the "Ports in Authorized and Unauthorized States" section on page 7-4.

The specific exchange of EAP frames depends on the authentication method being used. Figure 7-2 shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

Figure 7-2 Message Exchange

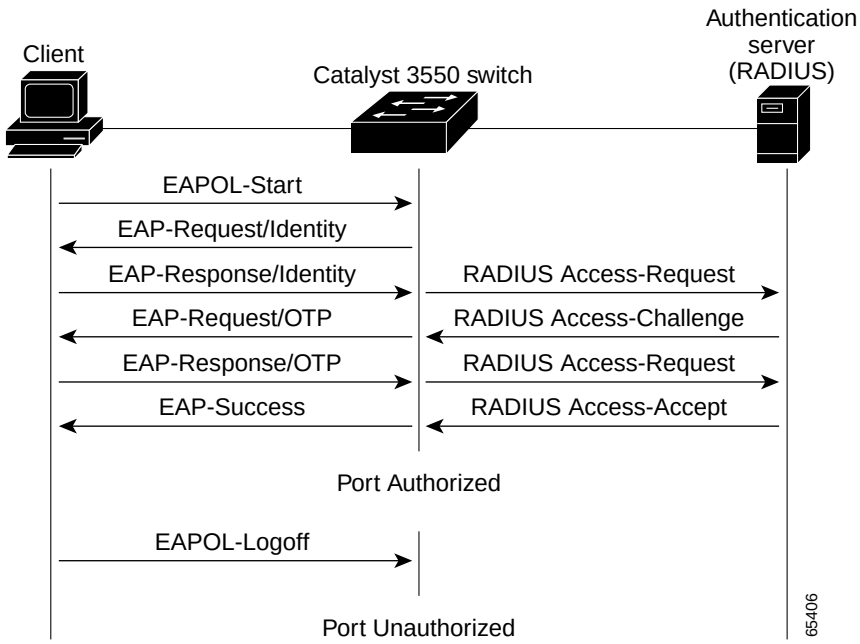

Ports in Authorized and Unauthorized States

The switch port state determines whether or not the client is granted access to the network. The port starts in the *unauthorized* state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the *authorized* state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1X is connected to an unauthorized 802.1X port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port that is not running the 802.1X protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized**—disables 802.1X authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized**—causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto**—enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client's MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

Supported Topologies

The 802.1X port-based authentication is supported in two topologies:

- Point-to-point
- Wireless LAN

In a point-to-point configuration (see Figure 7-1 on page 7-2), only one client can be connected to the 802.1X-enabled switch port. The switch detects the client when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

Figure 7-3 shows 802.1X port-based authentication in a wireless LAN. The 802.1X port is configured as a multiple-host port that becomes authorized as soon as one client is authenticated. When the port is authorized, all other hosts indirectly attached to the port are granted access to the network. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), the switch denies access to the network to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and the wireless access point acts as a client to the switch.

Figure 7-3 Wireless LAN Example

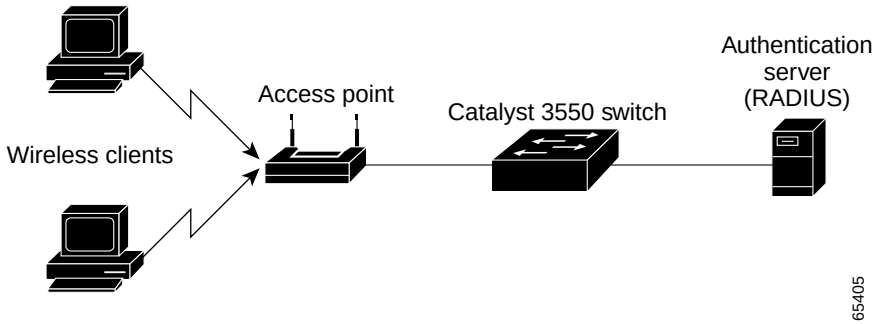

Configuring 802.1X Authentication

The section describes how to configure 802.1X port-based authentication on your switch. It contains this configuration information:

- Default 802.1X Configuration, page 7-6
- 802.1X Configuration Guidelines, page 7-7
- Enabling 802.1X Authentication, page 7-8 (required)
- Configuring the Switch-to-RADIUS-Server Communication, page 7-9 (required)
- Enabling Periodic Re-Authentication, page 7-10 (optional)
- Manually Re-Authenticating a Client Connected to a Port, page 7-11 (optional)
- Changing the Quiet Period, page 7-11 (optional)
- Changing the Switch-to-Client Retransmission Time, page 7-12 (optional)
- Setting the Switch-to-Client Frame-Retransmission Number, page 7-13 (optional)
- Enabling Multiple Hosts, page 7-13 (optional)
- Resetting the 802.1X Configuration to the Default Values, page 7-14 (optional)

Default 802.1X Configuration

Table 7-1 shows the default 802.1X configuration.

Table 7-1 Default 802.1X Configuration

Feature	Default Setting
Authentication, authorization, and accounting (AAA)	Disabled.
RADIUS server • IP address • UDP authentication port • Key	 • None specified. • 1812. • None specified.
Per-interface 802.1X protocol enable state	Disabled (force-authorized). The port transmits and receives normal traffic without 802.1X-based authentication of the client.
Periodic re-authentication	Disabled.
Number of seconds between re-authentication attempts	3600 seconds.
Quiet period	60 seconds (number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client).
Retransmission time	30 seconds (number of seconds that the switch should wait for a response to an EAP request/identity frame from the client before retransmitting the request).
Maximum retransmission number	2 times (number of times that the switch will send an EAP-request/identity frame before restarting the authentication process).
Multiple host support	Disabled.
Client timeout period	30 seconds (when relaying a request from the authentication server to the client, the amount of time the switch waits for a response before retransmitting the request to the client.
Authentication server timeout period	30 seconds (when relaying a response from the client to the authentication server, the amount of time the switch waits for a reply before retransmitting the response to the server. This setting is not configurable.)

802.1X Configuration Guidelines

These are the 802.1X authentication configuration guidelines:

- When 802.1X is enabled, ports are authenticated before any other Layer 2 or Layer 3 features are enabled.
- The 802.1X protocol is supported on both Layer 2 static-access ports and Layer 3 routed ports, but it is not supported on these port types:
  - Trunk port—If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.
  - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, the port mode is not changed.
  - Dynamic-access ports—If you try to enable 802.1X on a dynamic-access (VLAN Query Protocol [VQP]) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
  - EtherChannel port—Before enabling 802.1X on the port, you must first remove it from the EtherChannel. If you try to enable 802.1X on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1X is not enabled. If you enable 802.1X on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
  - Secure port—You cannot configure a secure port as an 802.1X port. If you try to enable 802.1X on a secure port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to a secure port, an error message appears, and the security settings are not changed.
  - Switch Port Analyzer (SPAN) destination port—You can enable 802.1X on a port that is a SPAN destination port; however, 802.1X is disabled until the port is removed as a SPAN destination. You can enable 802.1X on a SPAN source port.

Enabling 802.1X Authentication

To enable 802.1X port-based authentication, you must enable AAA and specify the authentication method list. A method list describes the sequence and authentication methods to be queried to authenticate a user.

The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure 802.1X port-based authentication. This procedure is required.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**aaa new-model**	Enable AAA.
Step 3	**aaa authentication dot1x** {**default**} *method1* [*method2*...]	Create an 802.1X authentication method list. To create a default list that is used when a named list is *not* specified in the **authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. Enter at least one of these keywords: • **group radius**—Use the list of all RADIUS servers for authentication. • **none**—Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.
Step 4	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be enabled for 802.1X authentication.
Step 5	**dot1x port-control auto**	Enable 802.1X authentication on the interface. For feature interaction information with trunk, dynamic, dynamic-access, EtherChannel, secure, and SPAN ports, see the "802.1X Configuration Guidelines" section on page 7-7.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show dot1x**	Verify your entries. Check the Status column in the 802.1X Port Summary section of the display. An *enabled* status means the port-control value is set either to **auto** or to **force-unauthorized**.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable 802.1X AAA authentication, use the **no aaa authentication dot1x** {**default** | *list-name*} *method1* [*method2*...] global configuration command. To disable 802.1X authentication, use the **dot1x port-control force-authorized** or the **no dot1x port-control** interface configuration command.

This example shows how to enable AAA and 802.1X on Fast Ethernet port 0/1:

```
Switch# configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa authentication dot1x default group radius
Switch(config)# interface fastethernet0/1
Switch(config-if)# dot1x port-control auto
Switch(config-if)# end
```

Configuring the Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by their host name or IP address, host name and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the fail-over backup to the first one. The RADIUS host entries are tried in the order that they were configured.

Beginning in privileged EXEC mode, follow these steps to configure the RADIUS server parameters on the switch. This procedure is required.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**radius-server host** {*hostname* | *ip-address*} **auth-port** *port-number* **key** *string*	Configure the RADIUS server parameters on the switch. For *hostname* | *ip-address*, specify the host name or IP address of the remote RADIUS server. For **auth-port** *port-number*, specify the UDP destination port for authentication requests. The default is 1812. For **key** *string*, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server. **Note** Always configure the key as the last item in the **radius-server host** command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon. If you want to use multiple RADIUS servers, re-enter this command.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete the specified RADIUS server, use the **no radius-server host** {*hostname* | *ip-address*} global configuration command.

This example shows how to specify the server with IP address 172.20.39.46 as the RADIUS server, to use port 1612 as the authorization port, and to set the encryption key to *rad123*, matching the key on the RADIUS server:

```
Switch(config)# radius-server host 172.l20.39.46 auth-port 1612 key rad123
```

You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using the **radius-server host** global configuration command. If you want to configure these options on a per-server basis, use the **radius-server timeout**, **radius-server retransmit**, and the **radius-server key** global configuration commands. For more information, see the "Configuring Settings for All RADIUS Servers" section on page 6-28.

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Enabling Periodic Re-Authentication

You can enable periodic 802.1X client re-authentication and specify how often it occurs. If you do not specify a time period before enabling re-authentication, the number of seconds between re-authentication attempts is 3600.

Automatic 802.1X client re-authentication is a global setting and cannot be set for clients connected to individual ports. To manually re-authenticate the client connected to a specific port, see the "Manually Re-Authenticating a Client Connected to a Port" section on page 7-11.

Beginning in privileged EXEC mode, follow these steps to enable periodic re-authentication of the client and to configure the number of seconds between re-authentication attempts:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**dot1x re-authentication**	Enable periodic re-authentication of the client, which is disabled by default.
Step 3	**dot1x timeout re-authperiod** *seconds*	Set the number of seconds between re-authentication attempts. The range is 1 to 4294967295; the default is 3600 seconds. This command affects the behavior of the switch only if periodic re-authentication is enabled.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show dot1x**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable periodic re-authentication, use the **no dot1x re-authentication** global configuration command.To return to the default number of seconds between re-authentication attempts, use the **no dot1x timeout re-authperiod** global configuration command.

This example shows how to enable periodic re-authentication and set the number of seconds between re-authentication attempts to 4000:

```
Switch(config)# dot1x re-authentication
Switch(config)# dot1x timeout re-authperiod 4000
```

Manually Re-Authenticating a Client Connected to a Port

You can manually re-authenticate the client connected to a specific port at any time by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command. If you want to enable or disable periodic re-authentication, see the "Enabling Periodic Re-Authentication" section on page 7-10.

This example shows how to manually re-authenticate the client connected to Fast Ethernet port 0/1:

```
Switch# dot1x re-authenticate interface fastethernet0/1
Starting reauthentication on FastEthernet0/1
```

Changing the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time, and then tries again. The idle time is determined by the quiet-period value. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a smaller number than the default.

Beginning in privileged EXEC mode, follow these steps to change the quiet period:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**dot1x timeout quiet-period** *seconds*	Set the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client. The range is 0 to 65535 seconds; the default is 60.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show dot1x**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default quiet time, use the **no dot1x timeout quiet-period** global configuration command.

This example shows how to set the quiet time on the switch to 30 seconds:

```
Switch(config)# dot1x timeout quiet-period 30
```

Changing the Switch-to-Client Retransmission Time

The client responds to the EAP-request/identity frame from the switch with an EAP-response/identity frame. If the switch does not receive this response, it waits a set period of time (known as the retransmission time) and then retransmits the frame.

Note You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to change the amount of time that the switch waits for client notification:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**dot1x timeout tx-period** *seconds*	Set the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before retransmitting the request. The range is 1 to 65535 seconds; the default is 30.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show dot1x**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default retransmission time, use the **no dot1x timeout tx-period** global configuration command.

This example shows how to set 60 as the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before retransmitting the request:

```
Switch(config)# dot1x timeout tx-period 60
```

Setting the Switch-to-Client Frame-Retransmission Number

In addition to changing the switch-to-client retransmission time, you can change the number of times that the switch sends an EAP-request/identity frame (assuming no response is received) to the client before restarting the authentication process.

Note You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to set the switch-to-client frame-retransmission number:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**dot1x max-req** *count*	Set the number of times that the switch sends an EAP-request/identity frame to the client before restarting the authentication process. The range is 1 to 10; the default is 2.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show dot1x**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default retransmission number, use the **no dot1x max-req** global configuration command.

This example shows how to set 5 as the number of times that the switch sends an EAP-request/identity request before restarting the authentication process:

```
Switch(config)# dot1x max-req 5
```

Enabling Multiple Hosts

You can attach multiple hosts to a single 802.1X-enabled port as shown in Figure 7-3 on page 7-5. In this mode, only one of the attached hosts must be successfully authorized for all hosts to be granted network access. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), all attached clients are denied access to the network.

Beginning in privileged EXEC mode, follow these steps to allow multiple hosts (clients) on an 802.1X-authorized port that has the **dot1x port-control** interface configuration command set to **auto**.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to which multiple hosts are indirectly attached.
Step 3	**dot1x multiple-hosts**	Allow multiple hosts (clients) on an 802.1X-authorized port. Make sure that the **dot1x port-control** interface configuration command set is set to **auto** for the specified interface.

	Command	Purpose
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show dot1x interface** *interface-id*	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable multiple hosts on the port, use the **no dot1x multiple-hosts** interface configuration command.

This example shows how to enable 802.1X on Fast Ethernet interface 0/1 and to allow multiple hosts:

```
Switch(config)# interface fastethernet0/1
Switch(config-if)# dot1x port-control auto
Switch(config-if)# dot1x multiple-hosts
```

Resetting the 802.1X Configuration to the Default Values

Beginning in privileged EXEC mode, follow these steps to reset the 802.1X configuration to the default values:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**dot1x default**	Reset the configurable 802.1X parameters to the default values.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show dot1x**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Displaying 802.1X Statistics and Status

To display 802.1X statistics for all interfaces, use the **show dot1x statistics** privileged EXEC command. To display 802.1X statistics for a specific interface, use the **show dot1x statistics interface** *interface-id* privileged EXEC command.

To display the 802.1X administrative and operational status for the switch, use the **show dot1x** privileged EXEC command. To display the 802.1X administrative and operational status for a specific interface, use the **show dot1x interface** *interface-id* privileged EXEC command.

For detailed information about the fields in these displays, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Configuring Interface Characteristics

This chapter defines the types of interfaces on the switch and describes how to configure them. The chapter has these sections:

- Understanding Interface Types, page 8-1
- Using the Interface Command, page 8-6
- Configuring Layer 2 Interfaces, page 8-12
- Monitoring and Maintaining the Layer 2 Interface, page 8-18
- Configuring Layer 3 Interfaces, page 8-22

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the online *Cisco IOS Interface Command Reference for Release 12.1.*

Understanding Interface Types

This section describe the different types of interfaces supported by the switch with references to chapters that contain more detailed information about configuring these interface types. The rest of the chapter describes configuration procedures for physical interface characteristics.

These sections are included:

- Port-Based VLANs, page 8-2
- Switch Ports, page 8-2
- EtherChannel Port Groups, page 8-3
- Switch Virtual Interfaces, page 8-4
- Routed Ports, page 8-4
- Connecting Interfaces, page 8-5

Port-Based VLANs

A VLAN is a switched network that is logically segmented by function, team, or application, without regard to the physical location of the users. For more information about VLANs, see Chapter 9, "Creating and Maintaining VLANs." Packets received on a port are forwarded only to ports that belong to the same VLAN as the receiving port. Network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs.

VLAN partitions provide hard firewalls for traffic in the VLAN, and each VLAN has its own MAC address table. A VLAN comes into existence when a local port is configured to be associated with the VLAN, when the VLAN Trunking Protocol (VTP) learns of its existence from a neighbor on a trunk, or when a user adds a VLAN to the local VTP database.

To configure VLANs, use the **vlan database** privileged EXEC command to enter VLAN configuration mode.

Add ports to a VLAN by using the **switchport** interface configuration commands:

- Identify the interface.
- For a trunk port, set trunk characteristics, and if desired, define the VLANs to which it can belong.
- For an access port, set and define the VLAN to which it belongs.

Switch Ports

Switch ports are Layer 2 only interfaces associated with a physical port. A switch port can be either an access port or a trunk port. You can configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to determine if a switch port should be an access port or a trunk port by negotiating with the port on the other end of the link. Switch ports are used for managing the physical interface and associated Layer 2 protocols and do not handle routing or bridging.

Configure switch ports (access ports and trunk ports) by using the **switchport** interface configuration commands. For detailed information about configuring access ports and trunk ports, see Chapter 9, "Creating and Maintaining VLANs."

Access Ports

An access port carries the traffic of and belongs to only one VLAN. Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port. If an access port receives a tagged packet (Inter-Switch Link [ISL] or 802.1Q tagged), the packet is dropped, the source address is not learned, and the frame is counted in the *No destination* statistic.

Two types of access ports are supported:

- Static access ports are manually assigned to a VLAN.
- VLAN membership of dynamic access ports is learned through incoming packets. By default, a dynamic access port is a member of no VLAN, and forwarding to and from the port is enabled only when the VLAN membership of the port is discovered. In the Catalyst 3550 switch, dynamic access ports are assigned to a VLAN by a VLAN Membership Policy Server (VMPS). The VMPS can be a Catalyst 6000 series switch; the Catalyst 3550 switch does not support the function of a VMPS.

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Two types of trunk ports are supported:

- In an ISL trunk port, all received packets are expected to be encapsulated with an ISL header, and all transmitted packets are sent with an ISL header. Native (non-tagged) frames received from an ISL trunk port are dropped.
- An IEEE 802.1Q trunk port supports simultaneous tagged and untagged traffic. An 802.1Q trunk port is assigned a default Port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a NULL VLAN ID are assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Although by default, a trunk port is a member of every VLAN known to the VTP, you can limit VLAN membership by configuring an allowed list of VLANs for each trunk port. The list of allowed VLANs does not affect any other port but the associated trunk port. By default, all possible VLANs (VLAN ID 1 to 1005) are in the allowed list. A trunk port can only become a member of a VLAN if VTP knows of the VLAN and the VLAN is in the enabled state. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of that VLAN and traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of a new, enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and no traffic for the VLAN is forwarded to or from the port.

Note VLAN 1 cannot be excluded from the allowed list.

For more information about trunk ports, see Chapter 9, “Creating and Maintaining VLANs.”

EtherChannel Port Groups

EtherChannel port groups provide the ability to treat multiple switch ports as one switch port. These port groups act as a single logical port for high-bandwidth connections between switches or between switches and servers. An EtherChannel balances the traffic load across the links in the channel. If a link within the EtherChannel fails, traffic previously carried over the failed link changes to the remaining links. You can group multiple trunk ports into one logical trunk port or group multiple access ports into one logical access port. Most protocols operate over either single ports or aggregated switch ports and do not recognize the physical ports within the port group. Exceptions are the DTP, the Cisco Discovery Protocol (CDP), and the Port Aggregation Protocol (PAgP), which operate only on physical ports.

When you configure an EtherChannel, you create a port-channel logical interface and assign an interface to the EtherChannel. For Layer 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command. For Layer 2 interfaces, the logical interface is dynamically created. For both Layer 3 and 2 interfaces, you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together. For more information, see Chapter 21, “Configuring EtherChannel.”

Switch Virtual Interfaces

A switch virtual interface (SVI) represents a VLAN of switch ports as one interface to the routing or bridging function in the system. Only one SVI can be associated with a VLAN, but you need to configure an SVI for a VLAN only when you wish to route between VLANs, fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. By default, an SVI is created for the default VLAN (VLAN 1) to permit remote switch administration. Additional SVIs must be explicitly configured. In Layer 2 mode, SVIs provide IP host connectivity only to the system; in Layer 3 mode, you can configure routing across SVIs.

Note To use SVIs in Layer 3 mode, you must have the enhanced multilayer software image (EMI) installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

SVIs are created the first time that you enter the **vlan** interface configuration command for a VLAN interface. The VLAN corresponds to the VLAN tag associated with data frames on an ISL or 802.1Q encapsulated trunk or the VLAN ID configured for an access port. Configure a VLAN interface for each VLAN for which you want to route traffic, and assign it an IP address. For more information, see the "Configuring IP Addressing" section on page 22-4.

SVIs support routing protocol and bridging configurations. For more information about configuring IP routing, see Chapter 22, "Configuring IP Unicast Routing," Chapter 24, "Configuring IP Multicast Routing,"and Chapter 26, "Configuring Fallback Bridging."

Routed Ports

A routed port is a physical port that acts like a port on a router; it does not have to be connected to a router. A routed port is not associated with a particular VLAN, as is an access port. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed ports can be configured with a Layer 3 routing protocol.

Note To configure routed ports, you must have the EMI installed on your switch.

Configure routed ports by putting the interface into Layer 3 mode with the **no switchport** interface configuration command. Then assign an IP address to the port, enable routing, and assign routing protocol characteristics by using the **ip routing** and **router** *protocol* global configuration commands.

Caution Entering a **no switchport** interface configuration command shuts the interface down and then re-enables it, which might generate messages on the device to which the interface is connected. Furthermore, when you use this command to put the interface into Layer 3 mode, you are deleting any Layer 2 characteristics configured on the interface.

The number of routed ports and SVIs that you can configure is not limited by software; however, the interrelationship between this number and the number of other features being configured might have an impact on CPU utilization because of hardware limitations. For more information about feature combinations, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

For more information about IP unicast and multicast routing and routing protocols, see Chapter 22, "Configuring IP Unicast Routing" and Chapter 24, "Configuring IP Multicast Routing."

Connecting Interfaces

Devices within a single VLAN can communicate directly through any switch. Ports in different VLANs cannot exchange data without going through a routing device or interface.

With a standard Layer 2 switch, ports in different VLANs have to exchange information through a router. In the configuration shown in Figure 8-1, when Host A in VLAN 20 sends data to Host B in VLAN 30, it must go from Host A to the switch, to the router, back to the switch, and then to Host B.

Figure 8-1 Connecting VLANs with Layer 2 Switches

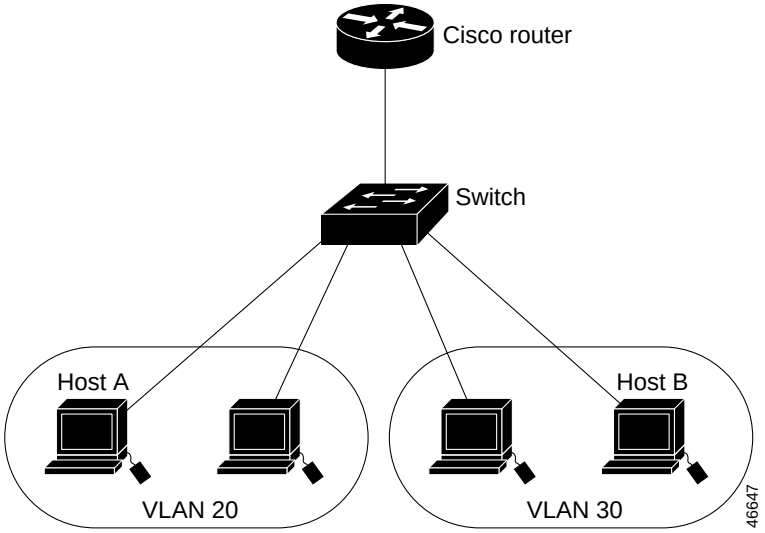

By using the Catalyst 3550 with the enhanced multilayer software image installed, when you configure VLAN 20 and VLAN 30 each with an SVI to which an IP address is assigned, packets can be sent from Host A to Host B directly through the Catalyst 3550 switch with no need for an external router (Figure 8-2).

Figure 8-2 Connecting VLANs with the Catalyst 3550 Multilayer Switch

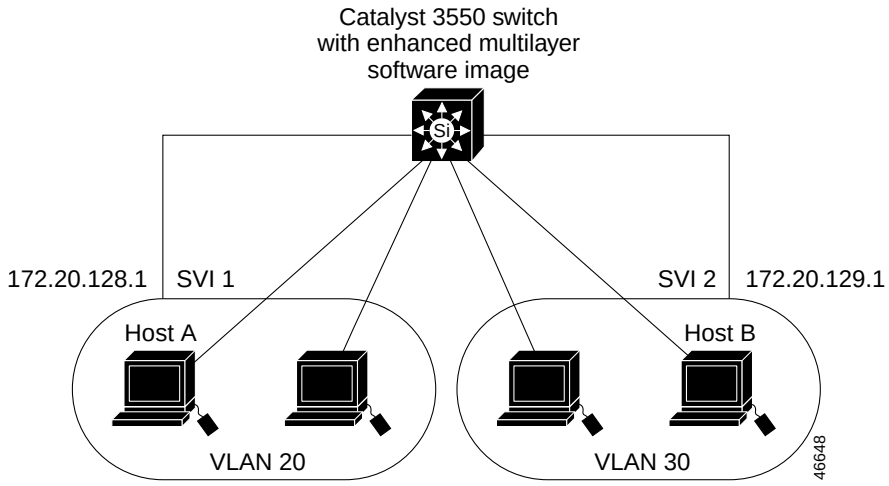

The Catalyst 3550 switch with the enhanced multilayer software image supports two methods of forwarding traffic between interfaces: routing and fallback bridging. Whenever possible, to maintain high performance, forwarding is done by switch hardware. However, only IP version 4 packets with Ethernet II encapsulation can be routed in hardware. All other types of traffic can be fallback bridged by hardware.

- The routing function can be enabled on all SVIs and routed ports. The Catalyst 3550 switches with the enhanced multilayer software image route only IP traffic. When IP routing protocol parameters and address configuration are added to an SVI or routed port, any IP traffic received from these ports is routed. For more information, see Chapter 22, "Configuring IP Unicast Routing," Chapter 24, "Configuring IP Multicast Routing," and Chapter 25, "Configuring MSDP."
- Fallback bridging forwards traffic that the switch with the enhanced multilayer software image does not route or traffic belonging to a nonroutable protocol, such as DECnet. Fallback bridging connects multiple VLANs into one bridge domain by bridging between two or more SVIs or routed ports. When configuring fallback bridging, you assign SVIs or routed ports to bridge groups with each SVI or routed port assigned to only one bridge group. All interfaces in the same group belong to the same bridge domain. For more information, see Chapter 26, "Configuring Fallback Bridging."

Using the Interface Command

The Catalyst 3550 switch supports these interface types:

- Physical ports—including switch ports and routed ports
- VLANs—Switch Virtual Interface
- Port-channels—EtherChannel of interfaces

You can also configure a range of interfaces (see the "Configuring a Range of Interfaces" section on page 8-9).

To configure a physical interface (port), enter interface configuration mode, and specify the interface type, slot, and number.

- Type—Fast Ethernet (fastethernet or fa) for 10/100 Ethernet or Gigabit Ethernet (gigabitethernet or gi)
- Slot—The slot number on the switch. On the Catalyst 3550 switch, the slot number is 0.
- Port number—The interface number on the switch. The port numbers always begin at 1, starting at the left when facing the front of the switch, for example, gigabitethernet 0/1, gigabitethernet 0/2. If there is more than one media type (for example, 10/100 ports and Gigabit Ethernet ports), the port number starts again with the second media: fastethernet0/1, fastethernet0/2.

You can identify physical interfaces by physically checking the interface location on the switch. You can also use the IOS **show** privileged EXEC commands to display information about a specific interface or all the interfaces on the switch. The remainder of this chapter primarily provides physical interface configuration procedures.

Procedures for Configuring Interfaces

These general instructions apply to all interface configuration processes.

Step 1 Enter the **configure terminal** command at the privileged EXEC prompt:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#
```

Step 2 Enter the **interface** global configuration command. Identify the interface type and the number of the connector. In this example, Gigabit Ethernet interface 0/1 is selected:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)#
```

Note You do not need to add a space between the interface type and interface number. For example, in the preceding line, you can specify either **gigabitethernet 0/1**, **gigabitethernet0/1**, **gi 0/1**, or **gi0/1**.

Step 3 Follow each **interface** command with the interface configuration commands your particular interface requires. The commands you enter define the protocols and applications that will run on the interface. The commands are collected and applied to the interface when you enter another interface command or enter **end** to return to privileged EXEC mode.

You can also configure a range of interfaces by using the **interface range** or **interface range macro** global configuration commands. Interfaces configured in a range must be the same type and must be configured with the same feature options.

Step 4 After you configure an interface, verify its status by using the **show** privileged EXEC commands listed in the "Monitoring and Maintaining the Layer 2 Interface" section on page 8-18.

Enter the **show interfaces** privileged EXEC command to see a list of all interfaces on or configured for the switch. A report is provided for each interface that the device supports or for the specified interface:

```
Switch# show interfaces
Vlan1 is up, line protocol is up
  Hardware is EtherSVI, address is 0000.0000.0000 (bia 0000.0000.00
  Internet address is 10.1.1.64/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:35, output 2d14h, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 1 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     264251 packets input, 163850228 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     380 packets output, 26796 bytes, 0 underruns
     0 output errors, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
FastEthernet0/1 is up, line protocol is down
  Hardware is Fast Ethernet, address is 0000.0000.0001 (bia 0000.00
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 2 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out

<output truncated>

GigabitEthernet0/2 is up, line protocol is down
  Hardware is Gigabit Ethernet, address is 0000.0000.001a (b
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
```

```
5 minute output rate 0 bits/sec, 0 packets/sec
   0 packets input, 0 bytes, 0 no buffer
   Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
   0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
   0 input packets with dribble condition detected
   0 packets output, 0 bytes, 0 underruns
   0 output errors, 0 collisions, 2 interface resets
   0 babbles, 0 late collision, 0 deferred
   0 lost carrier, 0 no carrier
   0 output buffer failures, 0 output buffers swapped out
```

Configuring a Range of Interfaces

You can use the **interface range** global configuration command to configure multiple interfaces with the same configuration parameters. When you enter the interface range configuration mode, all command parameters that you enter are attributed to all interfaces within that range until you exit this mode.

Beginning in privileged EXEC mode, follow these steps to configure a range of interfaces with the same parameters:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface range** {*port-range* | **macro** *macro_name*}	Enter interface range configuration mode by entering the range of interfaces (VLANs or physical ports) to be configured. • You can use the **interface range** command to configure up to five port ranges or a previously defined macro. • The **macro** variable is explained in the "Configuring and Using Interface Range Macros" section on page 8-11. • Each comma-separated *port-range* must consist of the same port type. You do not need to enter spaces before or after the comma. • When you define a range, the space between the first port and the hyphen is required.
Step 3		You can now use the normal configuration commands to apply the configuration parameters to all interfaces in the range.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show interfaces** [*interface-id*]	Verify the configuration of the interfaces in the range.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

When using the **interface range** global configuration command, note these guidelines:

- Valid entries for *port-range*:
  - **vlan** *vlan-ID - vlan-ID*
  - **fastethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **gigabitethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **port-channel** *port-channel-number - port-channel-number*, where *port-channel-number* is from 1 to 64

- You must add a space between the interface numbers and the hyphen when using the **interface range** command. For example, the command **interface range gigabitethernet 0/1 - 5** is a valid range; the command **interface range gigabitethernet 0/1-5** is not a valid range.
- The **interface range** command works only with VLAN interfaces that have been configured with the **interface vlan** command (the **show running-config** privileged EXEC command displays the configured VLAN interfaces). VLAN interfaces not displayed by the **show running-config** command cannot be used with the **interface range** command.
- All interfaces in a range must be the same type; that is, all Fast Ethernet ports, all Gigabit Ethernet ports, all EtherChannel ports, or all SVIs.

This example shows how to use the **interface range** global configuration command to enable Gigabit Ethernet interfaces 0/1 to 0/5:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/1 - 5
Switch(config-if-range)# no shutdown
Switch(config-if-range)#
*Oct  6 08:24:35: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
*Oct  6 08:24:35: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
*Oct  6 08:24:35: %LINK-3-UPDOWN: Interface GigabitEthernet0/3, changed state to up
*Oct  6 08:24:35: %LINK-3-UPDOWN: Interface GigabitEthernet0/4, changed state to up
*Oct  6 08:24:35: %LINK-3-UPDOWN: Interface GigabitEthernet0/5, changed state to up
*Oct  6 08:24:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/05,
changed state to up
*Oct  6 08:24:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/3,
changed state to up
*Oct  6 08:24:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/4,
changed state to up
```

This example shows how to use a comma to add different interface type strings to the range to enable all Gigabit Ethernet interfaces in the range 0/1 to 0/3 and both Gigabit Ethernet interfaces 0/7 and 0/8:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/1 - 3, gigabitethernet0/7 - 8
Switch(config-if-range)# no shutdown
Switch(config-if-range)#
*Oct  6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
*Oct  6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
*Oct  6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/3, changed state to up
*Oct  6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/7, changed state to up
*Oct  6 08:29:28: %LINK-3-UPDOWN: Interface GigabitEthernet0/8, changed state to up
*Oct  6 08:29:29: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/ 7,
changed state to up
*Oct  6 08:29:29: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/ 2,
changed state to up
*Oct  6 08:29:29: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/ 4,
changed state to up
```

If you enter multiple configuration commands while you are in interface range mode, each command is executed as it is entered. The commands are not batched together and executed after you exit interface range mode. If you exit interface range configuration mode while the commands are being executed, some commands might not be executed on all interfaces in the range. Wait until the command prompt reappears before exiting interface range configuration mode.

Configuring and Using Interface Range Macros

You can create an interface range macro to automatically select a range of interfaces for configuration. Before you can use the **macro** keyword in the **interface range macro** global configuration command string, you must use the **define interface-range** global configuration command to define the macro.

Beginning in privileged EXEC mode, follow these steps to define an interface range macro:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**define interface-range** *macro_name interface-range*	Define the interface-range macro, and save it in NVRAM. • The *macro_name* is a 32-character maximum character string. • A macro can contain up to five comma-separated interface ranges. You do not need to enter spaces before or after the comma. • Each *interface-range* must consist of the same port type.
Step 3	**interface range macro** *macro_name*	Select the interface range to be configured using the values saved in the interface-range macro called *macro_name.* You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config | include define**	Show the defined interface range macro configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no define interface-range** *macro_name* global configuration command to delete a macro.

When using the **define interface-range** global configuration command, note these guidelines:

- Valid entries for *interface-range*:
  - **vlan** *vlan-ID - vlan-ID*
  - **fastethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **gigabitethernet** slot/{*first port*} - {*last port*}, where slot is **0**
  - **port-channel** *port-channel-number - port-channel-number*, where *port-channel-number* is from 1 to 64.
- You must add a space between the interface numbers and the hyphen when entering an *interface-range*. For example, **gigabitethernet 0/1 - 5** is a valid range; **gigabitethernet 0/1-5** is not a valid range.
- The VLAN interfaces (SVIs) must have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used as *interface-ranges*.
- All interfaces in a range must be the same type; that is, all Fast Ethernet ports, all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs, but you can combine multiple interface types in a macro.

This example shows how to define an interface-range macro named *enet_list* to select Gigabit Ethernet ports 1 to 4 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list gigabitethernet0/1 - 4
Switch(config)# end
Switch# show running-config | include define
define interface-range enet_list GigabitEthernet0/1 - 4
Switch#
```

This example shows how to create a multiple-interface macro named *macro1*:

```
Switch# configure terminal
Switch(config)# define interface-range macro1 gigabitethernet0/1 - 2, fastethernet0/5 - 7
Switch(config)# end
Switch#
```

This example shows how to enter interface range configuration mode for the interface-range macro *enet_list*:

```
Switch# configure terminal
Switch(config)# interface range macro enet_list
Switch(config-if-range)#
```

This example shows how to delete the interface-range macro *enet_list* and to verify that it has been deleted.

```
Switch# configure terminal
Switch(config)# no define interface-range enet_list
Switch# show run | include define
Switch#
```

Configuring Layer 2 Interfaces

These sections describe the default interface configuration and the optional features that you can configure on most physical interfaces:

- Default Layer 2 Ethernet Interface Configuration, page 8-13
- Configuring Interface Speed and Duplex Mode, page 8-14
- Configuring IEEE 802.3X Flow Control, page 8-16
- Adding a Description for an Interface, page 8-17

Caution If the interface is in Layer 3 mode, after entering interface configuration mode, you must enter the **switchport** interface configuration command without any parameters to put the interface into Layer 2 mode. This shuts down the interface and then re-enables it, which might generate messages on the device to which the interface is connected. Furthermore, when you use this command to put the interface into Layer 2 mode, you are deleting any Layer 3 characteristics configured on the interface.

Default Layer 2 Ethernet Interface Configuration

Table 8-1 shows the Layer 2 Ethernet interface default configuration. For more details on the VLAN parameters listed in the table, see Chapter 9, "Creating and Maintaining VLANs." For details on controlling traffic to the port, see Chapter 12, "Configuring Port-Based Traffic Control."

Table 8-1 Default Layer 2 Ethernet Interface Configuration

Feature	Default Setting
Operating mode	Layer 2 or switching mode (**switchport** command). **Note** Routing mode (**no switchport** command) is only an option when the enhanced multilayer software image is installed. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either SMI or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.
Allowed VLAN range	VLANs 1 – 1005.
Default VLAN (for access ports)	VLAN 1.
Native VLAN (for 802.1Q trunks)	VLAN 1.
VLAN trunking	Switchport mode dynamic desirable (supports DTP).
Port enable state	All ports are enabled.
Port description	None defined.
Speed	Autonegotiate.
Duplex mode	Autonegotiate.
Flow control	Flow control set to *off* for **receive** and *desired* for **send** for 10/100/1000 Mb/s. For 10/100 Mb/s ports, **send** is always *off*.
EtherChannel (PAgP)	Disabled on all Ethernet ports. See Chapter 21, "Configuring EtherChannel."
Port blocking (unknown multicast and unknown unicast traffic)	Disabled (not blocked). See the "Configuring Port Blocking" section on page 12-6.
Broadcast, multicast, and unicast storm control	Disabled. See the "Default Storm Control Configuration" section on page 12-3.
Protected port	Disabled. See the "Configuring Protected Ports" section on page 12-5.
Port security	Disabled. See the "Default Port Security Configuration" section on page 12-9.
Port Fast	Disabled.

Configuring Interface Speed and Duplex Mode

These sections describe how to configure the interface speed and duplex mode:

- Configuration Guidelines, page 8-14
- Setting the Interface Speed and Duplex Parameters, page 8-14

Configuration Guidelines

When configuring an interface speed and duplex mode, note these guidelines:

- If both ends of the line support autonegotiation, we highly recommend the default autonegotiation settings.
- If one interface supports autonegotiation and the other end does not, configure duplex and speed on both interfaces; do not use the **auto** setting on the supported side.

Note GigaStack-to-GigaStack cascade connections operate in half-duplex mode, and GigaStack-to-GigaStack point-to-point connections operate in full-duplex mode.

Caution Changing the interface speed and duplex mode configuration might shut down and re-enable the interface during the reconfiguration.

Setting the Interface Speed and Duplex Parameters

You can configure interface speed on Fast Ethernet (10/100-Mbps) and Gigabit Ethernet (10/100/1000-Mbps) interfaces; you cannot configure speed on Gigabit Interface Converter (GBIC) interfaces. You can configure duplex mode on any Fast Ethernet or Gigabit Ethernet interfaces that are not set to autonegotiate; you cannot configure duplex mode on GBIC interfaces. Because collisions are major constrictions in Ethernet networks, full-duplex communication is an effective solution. Normally, 10-Mbps ports operate in half-duplex mode, which means that stations can either receive or send. In full-duplex mode, two stations can send and receive at the same time. When packets can flow in both directions simultaneously, effective Ethernet bandwidth doubles to 20 Mbps for 10-Mbps interfaces, to 200 Mbps for Fast Ethernet interfaces, and to 2 Gbps for Gigabit interfaces.

Note You cannot configure speed or duplex mode on Gigabit Interface Converter (GBIC) ports, but for certain types of GBICs, you can configure speed to not negotiate (**nonegotiate**) if connected to a device that does not support autonegotiation.

Beginning in privileged EXEC mode, follow these steps to set the speed and duplex mode for a physical interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode and the physical interface identification.
Step 3	**speed** {**10** | **100** | **1000** | **auto** | **nonegotiate**}	Enter the appropriate speed parameter for the interface, or enter **auto** or **nonegotiate**. **Note** The **1000** keyword is available only for 10/100/1000 Mbps ports. The **nonegotiate** keyword is available only for 1000BASE-SX, -LX, and -ZX GBIC ports.
Step 4	**duplex** {**auto** | **full** | **half**}	Enter the duplex parameter for the interface.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show interfaces** *interface-id*	Display the interface speed and duplex mode configuration.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no speed** and **no duplex** interface configuration commands to return the interface to the default speed and duplex settings (autonegotiate). To return all interface settings to the defaults, use the **default interface** *interface-id* interface configuration command.

This example shows how to set the interface speed to 10 Mbps and the duplex mode to half on FastEthernet interface 0/3 and to verify the configuration:

```
Switch# configure terminal
Switch(config)# interface fastethernet0/3
Switch(config-if)# speed 10
Switch(config-if)# duplex half
Switch(config-if)# end
Switch# show interfaces fastethernet0/3
FastEthernet0/3 is up, line protocol is down
  Hardware is Fast Ethernet, address is 0000.0000.0003 (bia 0000.0000.0003)
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Half-duplex, 10Mb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 2 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

Configuring IEEE 802.3X Flow Control

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If one port experiences congestion and cannot receive any more traffic, it notifies the other port to stop sending until the condition clears. When the local device detects any congestion at its end, it can notify the link partner or the remote device of the congestion by sending a pause frame. Upon receipt of a pause frame, the remote device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note You must not configure both IEEE 802.3X flowcontrol and quality of service (QoS) on a switch. Before configuring flowcontrol on an interface, use the **no mls qos** global configuration command to disable QoS on the switch.

Flow control can be implemented in two forms, symmetric and asymmetric. The symmetric implementation is suitable for point-to-point links, and asymmetric is suitable for hub-to-end node connections, where it is desirable for the hub to pause the end system, but not vice-versa. You use the **flowcontrol** interface configuration command to set the interface's ability to **receive** and **send** pause frames to **on**, **off**, or **desired**. The default state for Gigabit Ethernet ports is **receive off** and **send desired**. The default state for Fast Ethernet ports is **receive off** and **send off**.

Note On Catalyst 3550 switches, Gigabit Ethernet ports are capable of receiving and sending pause frames; Fast Ethernet ports can only receive pause frames. Therefore, for Fast Ethernet ports, only the conditions described with **send off** are applicable.

These rules apply to flow control settings on the device:

- **receive on** (or **desired**) and **send on**: Flow control operates in both directions; both the local and the remote devices can send pause frames to show link congestion.
- **receive on** (or **desired**) and **send desired**: The port can receive pause frames and can send pause frames if the attached device supports flow control.
- **receive on** (or **desired**) and **send off**: The port cannot send pause frames but can operate with an attached device that is required to or can send pause frames; the port can receive pause frames.
- **receive off** and **send on**: The port sends pause frames if the remote device supports flow control but cannot receive pause frames from the remote device.
- **receive off** and **send desired**: The port cannot receive pause frames but can send pause frames if the attached device supports flow control.
- **receive off** and **send off**: Flow control does not operate in either direction. In case of congestion, no indication is given to the link partner, and no pause frames are sent or received by either device.

Note For details on the command settings and the resulting flow control resolution on local and remote ports, refer to the **flowcontrol** interface configuration command in the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Beginning in privileged EXEC mode, follow these steps to configure flow control on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**no mls qos**	Disable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode and the physical interface to be configured.
Step 4	**flowcontrol** {**receive** | **send**} {**on** | **off** | **desired**}	Configure the flow control mode for the port. **Note** The **send** keyword is not available for 10/100 Mbps ports.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show interfaces** *interface-id*	Verify the interface flow control settings.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable flow control, use the **flowcontrol receive off** and **flowcontrol send off** interface configuration commands.

This example shows how to turn off all flow control on Gigabit Ethernet interface 0/1 and to display the results:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# flowcontrol receive off
Switch(config-if)# flowcontrol send off
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/1
GigabitEthernet0/1 is up, line protocol is down
  Hardware is Gigabit Ethernet, address is 0002.4b29.2e01 (bia 0002.4b29.2e01)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed
  input flow-control is off, output flow-control is off
<output truncated>
```

Adding a Description for an Interface

You can add a description about an interface to help you remember its function. The description appears in the output of these commands: **show configuration**, **show running-config**, and **show interfaces**.

Beginning in privileged EXEC mode, follow these steps to add a description for an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the interface for which you are adding a description.
Step 3	**description** *string*	Add a description (up to 240 characters) for an interface.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show interfaces** *interface-id* **description** or **show running-config**	Verify your entry.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no description** interface configuration command to delete the description.

This example shows how to add a description on Fast Ethernet interface 0/4 and to verify the description:

```
Switch# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface fastethernet0/4
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces fastethernet0/4 description
Interface Status         Protocol Description
Fa0/4     up             down     Connects to Marketing
```

Monitoring and Maintaining the Layer 2 Interface

You can perform the tasks in these sections to monitor and maintain the interfaces:

- Monitoring Interface and Controller Status, page 8-18
- Clearing and Resetting Interfaces and Counters, page 8-20
- Shutting Down and Restarting the Interface, page 8-21

Monitoring Interface and Controller Status

Commands entered at the privileged EXEC prompt display information about the interface, including the version of the software and the hardware, the controller status, and statistics about the interfaces. Table 8-2 lists some of these interface monitoring commands. (You can display the full list of **show** commands by using the **show ?** command at the privileged EXEC prompt.) These commands are fully described in the *Cisco IOS Interface Command Reference for Release 12.1*.

Table 8-2 Show Commands for Interfaces

Command	Purpose
show interfaces [*interface-id*]	Display the status and configuration of all interfaces or a specific interface.
show interfaces *interface-id* **status** [**err-disabled**]	Display interface status or a list of interfaces in error-disabled state.
show interfaces [*interface-id*] **switchport**	Display administrative and operational status of switching (nonrouting) ports. You can use this command to determine if a port is in routing or switching mode.
show interfaces [*interface-id*] **description**	Display the description configured on an interface or all interfaces and the interface status.

Table 8-2 Show Commands for Interfaces (continued)

Command	Purpose
show running-config	Display the running configuration in RAM.
show version	Display the hardware configuration, software version, the names and sources of configuration files, and the boot images.

This example shows how to display the status and configuration of Gigabit Ethernet interface 0/2:

```
Switch# show interfaces gigabitethernet0/2
GigabitEthernet0/2 is up, line protocol is up
  Hardware is Gigabit Ethernet, address is 0002.4b29.4400 (bia 0002.4b29.4400)
  Internet address is 192.20.135.21/24
  MTU 1500 bytes, BW 100000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:01, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     89245 packets input, 8451658 bytes, 0 no buffer
     Received 81551 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     60387 packets output, 5984015 bytes, 0 underruns
     0 output errors, 0 collisions, 16 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

This example shows how to display the status of all interfaces:

```
Switch# show interfaces status

Port      Name     Status          Vlan     Duplex     Speed    Type
Gi0/1     CubeA    connected       1        a-full     a-100    10/100/1000BaseTX
Gi0/2     CubeE    notconnect      1        auto       auto     10/100/1000BaseTX
Gi0/3     CubeF    disabled        1        auto       auto     10/100/1000BaseTX
Gi0/4     CubeG    notconnect      1        auto       auto     10/100/1000BaseTX
Gi0/5     CubeH    notconnect      routed   auto       auto     10/100/1000BaseTX
Gi0/6     CubeI    notconnect      routed   auto       auto     10/100/1000BaseTX
Gi0/7     CubeJ    connected       1        a-full     a-100    10/100/1000BaseTX
Gi0/8     CubeK    notconnect      1        auto       auto     10/100/1000BaseTX
Gi0/9     CubeL    disabled        1        auto       auto     10/100/1000BaseTX
Gi0/10    CubeB    notconnect      routed   auto       auto     10/100/1000BaseTX
Gi0/11    CubeC    notconnect      1        auto       auto     unknown
Gi0/12    CubeD    notconnect      1        auto       auto     unknown
```

This example shows how to display the status of switching ports:

```
Switch# show interfaces switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: True
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled

Name: Gi0/2
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: True
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled

<output truncated>
```

Clearing and Resetting Interfaces and Counters

Table 8-3 lists the privileged EXEC mode **clear** commands that you can use to clear counters and reset interfaces.

Table 8-3 Clear Commands for Interfaces

Command	Purpose
clear counters [*interface-id*]	Clear interface counters.
clear interface *interface-id*	Reset the hardware logic on an interface.
clear line [*number* | **console 0** | **vty** *number*]	Reset the hardware logic on an asynchronous serial line.

To clear the interface counters shown by the **show interfaces** privileged EXEC command, use the **clear counters** privileged EXEC command. The **clear counters** command clears all current interface counters from the interface unless optional arguments are specified to clear only a specific interface type from a specific interface number.

Note The **clear counters** privileged EXEC command does not clear counters retrieved by using Simple Network Management Protocol (SNMP), but only those seen with the **show interface** privileged EXEC command.

This example shows how to clear and reset the counters on Gigabit Ethernet interface 0/5:

```
Switch# clear counters gigabitethernet0/5
Clear "show interface" counters on this interface [confirm] y
Switch#
*Sep 30 08:42:55: %CLEAR-5-COUNTERS: Clear counter on interface GigabitEthernet0/5
by vty1 (171.69.115.10)
```

Use the **clear interface** or **clear line** privileged EXEC command to clear and reset an interface or serial line. Under most circumstances, you do not need to clear the hardware logic on interfaces or serial lines.

This example shows how to clear and reset Gigabit Ethernet interface 0/5:

```
Switch# clear interface gigabitethernet0/5
```

Shutting Down and Restarting the Interface

Shutting down an interface disables all functions on the specified interface and marks the interface as unavailable on all monitoring command displays. This information is communicated to other network servers through all dynamic routing protocols. The interface is not mentioned in any routing updates.

Beginning in privileged EXEC mode, follow these steps to shut down an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** {**vlan** *vlan-id*} | {{**fastethernet** | **gigabitethernet**} *interface-id*} | {**port-channel** *port-channel-number*}	Select the interface to be configured.
Step 3	**shutdown**	Shut down an interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entry.

Use the **no shutdown** interface configuration command to restart the interface.

This example shows how to shut down Gigabit Ethernet interface 0/5:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/5
Switch(config-if)# shutdown
Switch(config-if)#
*Sep 30 08:33:47: %LINK-5-CHANGED: Interface GigabitEthernet0/5, changed state to a
administratively down
```

This example shows how to re-enable Gigabit Ethernet interface 0/5:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/5
Switch(config-if)# no shutdown
Switch(config-if)#
*Sep 30 08:36:00: %LINK-3-UPDOWN: Interface GiogabitEthernet0/5, changed state to up
```

To verify that an interface is disabled, enter the **show interfaces** privileged EXEC command. A disabled interface is shown as *administratively down* in the **show interface** command display as with Gigabit Ethernet interface 0/1 in this example.

```
Switch# show interfaces
<output truncated>

GigabitEthernet0/2 is administratively down, line protocol is down
  Hardware is Gigabit Ethernet, address is 0002.4b29.4403 (bia 0002.4b29.4403)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed

<output truncated>
```

Configuring Layer 3 Interfaces

The Catalyst 3550 with the enhanced multilayer software image supports three types of Layer 3 interfaces for routing and bridging:

- SVIs: You should configure SVIs for any VLANs for which you want to route traffic. SVIs are created when you enter a VLAN ID following the **interface vlan** global configuration command. To delete an SVI, use the **no interface vlan** global configuration command.

  For information about assigning Layer 2 ports to VLANs, see Chapter 9, "Creating and Maintaining VLANs."

- Layer 3 EtherChannel ports: EtherChannel interfaces made up of routed ports.

  EtherChannel port interfaces are described in Chapter 21, "Configuring EtherChannel."

- Routed ports: Routed ports are physical ports configured to be in Layer 3 mode by using the **no switchport** interface configuration command.

Note A Layer 3 switch can have an IP address assigned to each routed port and SVI. The number of routed ports and SVIs that you can configure is not limited by software; however, the interrelationship between this number and the number of other features being configured might have an impact on CPU utilization because of hardware limitations. For more information about feature combinations, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

All Layer 3 interfaces require an IP address to route traffic. The following procedure shows how to configure an interface as a Layer 3 interface and how to assign an IP addresses to an interface.

Note If the physical port is in Layer 2 mode (the default), you must enter the **no switchport** interface configuration command to put the interface into Layer 3 mode. Entering a **no switchport** command disables and then re-enables the interface, which might generate messages on the device to which the interface is connected. When you use this command to put the interface into Layer 3 mode, you are also deleting any Layer 2 characteristics configured on the interface.

Beginning in privileged EXEC mode, follow these steps to configure a Layer 3 interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** {{**fastethernet** | **gigabitethernet**} *interface-id*} | {**vlan** *vlan-id*} | {**port-channel** *port-channel-number*}	Enter interface configuration mode, and enter the interface to be configured as a Layer 3 interface.
Step 3	**no switchport**	For physical ports only, enter Layer 3 mode.
Step 4	**ip address** *ip_address subnet_mask*	Configure the IP address and IP subnet.
Step 5	**no shutdown**	Enable the interface.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show interfaces** [*interface-id*] **show ip interface** [*interface-id*] **show running-config interface** [*interface-id*]	Verify the configuration.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove an IP address from an interface, use the **no ip address** interface configuration command.

This example shows how to configure an interface as a routed port and to assign it an IP address:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 192.20.135.21 255.255.255.0
Switch(config-if)# no shutdown
Switch(config-if)# end
```

This is an example of output from the **show interfaces** privileged EXEC command for an interface:

```
Switch(config)# show interfaces gigabitethernet0/2
GigabitEthernet0/2 is up, line protocol is up
  Hardware is Gigabit Ethernet, address is 0002.4b29.4400 (bia 0002.4b29.4400)
  Internet address is 192.20.135.21/24
  MTU 1500 bytes, BW 100000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:02, output 00:00:08, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     89604 packets input, 8480109 bytes, 0 no buffer
     Received 81848 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     60665 packets output, 6029820 bytes, 0 underruns
     0 output errors, 0 collisions, 16 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

This is an example of output from the **show ip interface** privileged EXEC command for an interface:

```
Switch# show ip interface gigabitethernet0/2
GigabitEthernet0/2 is up, line protocol is up
  Internet address is 192.20.135.21/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast, CEF
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect exclude is disabled
  BGP Policy Mapping is disabled
```

This is an example of output for the the **show running-config** privileged EXEC command for an interface:

```
Switch# show running-config interface gigabitethernet0/2
Building configuration...

Current configuration : 122 bytes
!
interface GigabitEthernet0/2
 no switchport
 ip address 192.20.135.21 255.255.255.0
 speed 100
 mls qos trust dscp
end
```

Creating and Maintaining VLANs

This chapter describes how to create and maintain VLANs. It includes information about VLAN modes, the VLAN Trunking Protocol (VTP) database, and the VLAN Membership Policy Server (VMPS).

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

The chapter includes these sections:

- Understanding VLANs, page 9-1
- Using the VLAN Trunking Protocol, page 9-3
- VLANs in the VTP Database, page 9-15
- Understanding VLAN Trunks, page 9-22
- Understanding VMPS, page 9-33

Understanding VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge as shown in Figure 9-1. Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of the Spanning Tree Protocol (STP).

Note Before you create VLANs, you must decide whether to use VTP to maintain global VLAN configuration for your network. For more information on VTP, see the "Using the VLAN Trunking Protocol" section on page 9-3.

Figure 9-1 shows an example of VLANs segmented into logically defined networks.

Figure 9-1 VLANs as Logically Defined Networks

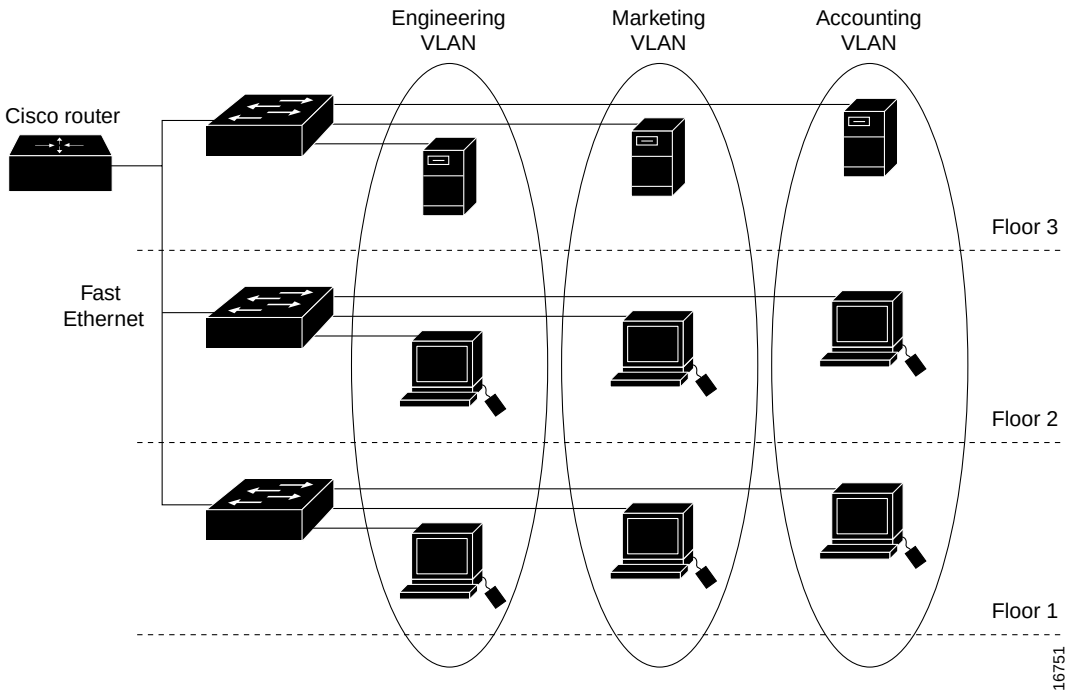

VLANs are often associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Interface VLAN membership on the switch is assigned manually on an interface-by-interface basis. When you assign switch interfaces to VLANs by using this method, it is known as interface-based, or static, VLAN membership.

Traffic between VLANs must be routed. A Catalyst 3550 switch with the enhanced mutilayer software image installed can route traffic between VLANs by using switch virtual interfaces (SVIs). An SVI must be explicitly configured and assigned an IP address to route traffic between VLANs. For more information, see the "Switch Virtual Interfaces" section on page 8-4 and the "Configuring Layer 3 Interfaces" section on page 8-22.

Number of Supported VLANs

The Catalyst 3550 switch supports 1005 VLANs in VTP client, server, and transparent modes. VLANs are identified with a number between 1 and 1001. VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs. The switch supports per-VLAN spanning tree (PVST) with a maximum of 128 spanning-tree instances. One spanning-tree instance is allowed per VLAN.

The switch supports both Inter-Switch Link (ISL) and IEEE 802.1Q trunking methods for sending VLAN traffic over Ethernet ports.

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that determines the kind of traffic the port carries and the number of VLANs to which it can belong. Table 9-1 lists the membership modes and characteristics.

Table 9-1 Port Membership Modes

Membership Mode	VLAN Membership Characteristics
Static-access	A static-access port can belong to one VLAN and is manually assigned by using the **switchport mode access** interface configuration command. For more information, see the "Assigning Static-Access Ports to a VLAN" section on page 9-19.
Trunk (ISL or IEEE 802.1Q)	A trunk is a member of all VLANs in the VLAN database by default, but membership can be limited by configuring the allowed-VLAN list. You can also modify the pruning-eligible list to block flooded traffic to VLANs on trunk ports that are included in the list. VTP maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP exchanges VLAN configuration messages with other switches over trunk links. Configure VLAN trunks using the **switchport mode trunk** interface configuration command. For more information, see the "Configuring an Ethernet Interface as a Trunk Port" section on page 9-25.
Dynamic access	A dynamic-access port can belong to one VLAN and is dynamically assigned by a VMPS. The VMPS can be a Catalyst 5000 or Catalyst 6000 series switch, for example, but never a Catalyst 3550 switch. You begin configuration by using the **switchport mode access** interface configuration command. For more information, see the "Configuring an Interface as a Layer 2 Dynamic Access Port" section on page 9-37.

For more detailed definitions of the modes and their functions, see Table 9-6 on page 9-23.

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis. For more information, see the "Managing the MAC Address Table" section on page 6-51.

Using the VLAN Trunking Protocol

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches.

The VTP Domain and VTP Modes

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches under the same administrative responsibility sharing the same VTP domain name. A switch can be in only one VTP domain.You make global VLAN configuration changes for the domain by using the command-line interface (CLI), Cluster Management Suite (CMS) software, or Simple Network Management Protocol (SNMP).

You can configure a supported switch to be in one of the VTP modes listed in Table 9-2.

Table 9-2 VTP Modes

VTP Mode	Description
VTP server	In this mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as the VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links. In VTP server mode, VLAN configurations are saved in nonvolatile RAM (NVRAM). VTP server is the default mode.
VTP client	A VTP client behaves like a VTP server, but you cannot create, change, or delete VLANs on a VTP client. In VTP client mode, VLAN configurations are not saved in NVRAM.
VTP transparent	VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive from other switches from their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode. In VTP transparent mode, VLAN configurations are saved in NVRAM, but they are not advertised to other switches.

By default, the switch is in VTP server mode and in the no-management-domain state until it receives an advertisement for a domain over a trunk link (a link that carries the traffic of multiple VLANs) or until you configure a domain name. Until the management domain name is specified or learned, you cannot create or modify VLANs on a VTP server, and VLAN information is not propagated over the network.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch then ignores advertisements with a different domain name or an earlier configuration revision number.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP advertisements are sent over all trunk connections, including Inter-Switch Link (ISL) and IEEE 802.1Q.

VTP maps VLANs dynamically across multiple LAN types with unique names and internal index associates. Mapping eliminates excessive device administration required from network administrators.

If you configure a switch for VTP transparent mode, you can create and modify VLANs, but the changes are not sent to other switches in the domain, and they affect only the individual switch.

The "Configuring VTP" section on page 9-8 provides tips and caveats for configuring VTP.

VTP Advertisements

Each switch in the VTP domain sends periodic global configuration advertisements from each trunk port to a reserved multicast address. Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.

Note Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements. For more information on trunk ports, see the "Understanding VLAN Trunks" section on page 9-22.

VTP advertisements distribute this global domain information:

- VTP domain name
- VTP configuration revision number
- Update identity and update timestamp
- MD5 digest VLAN configuration, including maximum transmission unit (MTU) size for each VLAN.
- Frame format

VTP advertisements distribute this VLAN information for each configured VLAN:

- VLAN IDs (ISL and 802.1Q)
- VLAN name
- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

VTP Version 2

If you use VTP in your network, you must decide whether to use version 1 or version 2.

VTP version 2 supports these features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring Bridge Relay Function (TrBRF) and Token Ring Concentrator Relay Function (TrCRF) VLANs. For more information about Token Ring VLANs, see the "VLANs in the VTP Database" section on page 9-15.
- Unrecognized Type-Length-Value (TLV) support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in NVRAM when the switch is operating in VTP server mode.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Because VTP version 2 supports only one domain, it forwards VTP messages in transparent mode without inspecting the version and domain name.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI, the Cluster Management Software (CMS), or SNMP. Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the MD5 digest on a received VTP message is correct, its information is accepted.

VTP Pruning

VTP pruning increases network available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to reach the destination devices. Without VTP pruning, a switch floods broadcast, multicast, and unknown unicast traffic across all trunk links within a VTP domain even though receiving switches might discard them. VTP pruning is disabled by default.

VTP pruning blocks unneeded flooded traffic to VLANs on trunk ports that are included in the pruning-eligible list. Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on Catalyst 3550 trunk ports. If the VLANs are configured as pruning-ineligible, the flooding continues. VTP pruning is supported with VTP version 1 and version 2.

Figure 9-2 shows a switched network without VTP pruning enabled. Port 1 on Switch 1 and Port 2 on Switch 4 are assigned to the Red VLAN. If a broadcast is sent from the host connected to Switch 1, Switch 1 floods the broadcast and every switch in the network receives it, even though Switches 3, 5, and 6 have no ports in the Red VLAN.

Figure 9-2 Flooding Traffic without VTP Pruning

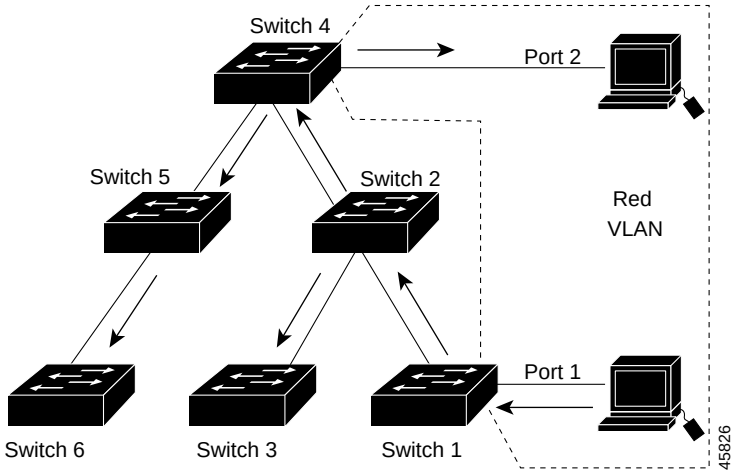

Figure 9-3 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch 1 is not forwarded to Switches 3, 5, and 6 because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch 2 and Port 4 on Switch 4).

Figure 9-3 Optimized Flooded Traffic with VTP Pruning

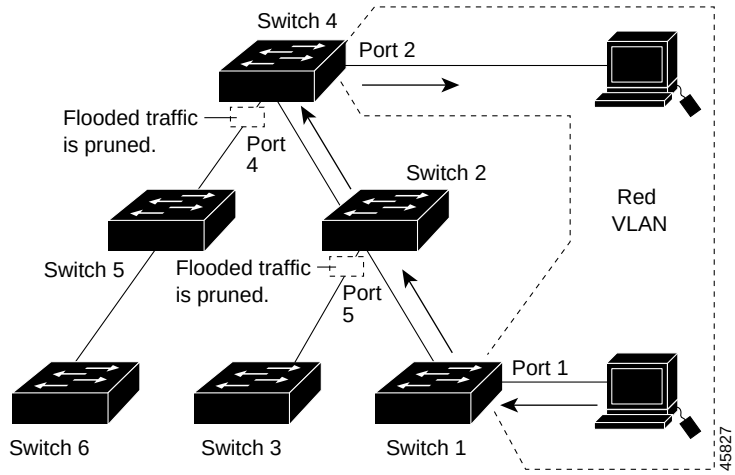

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. See the "Enabling VTP Pruning" section on page 9-13. VTP pruning takes effect several seconds after you enable it. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 is always pruning-ineligible; traffic from VLAN 1 cannot be pruned.

VTP pruning is not designed to function in VTP transparent mode. If one or more switches in the network are in VTP transparent mode, you should do one of these:

- Turn off VTP pruning in the entire network.
- Turn off VTP pruning by making all VLANs on the trunk of the switch upstream to the VTP transparent switch pruning ineligible.

To configure VTP pruning on an interface, use the **switchport trunk pruning vlan** interface configuration command (see the "Changing the Pruning-Eligible List" section on page 9-28). VTP pruning operates when an interface is trunking. You can set VLAN pruning-eligibility, whether or not VTP pruning is enabled for the VTP domain, whether or not any given VLAN exists, and whether or not the interface is currently trunking.

Configuring VTP

This section includes procedures for configuring VTP. These sections are included:

- Default VTP Configuration, page 9-8
- VTP Configuration Guidelines, page 9-8
- Configuring a VTP Server, page 9-10
- Configuring a VTP Client, page 9-11
- Disabling VTP (VTP Transparent Mode), page 9-11
- Enabling VTP Version 2, page 9-12
- Enabling VTP Pruning, page 9-13
- Monitoring VTP, page 9-13

Default VTP Configuration

Table 9-3 shows the default VTP configuration.

Table 9-3 Default VTP Configuration

Feature	Default Setting
VTP domain name	Null.
VTP mode	Server.
VTP version 2 enable state	Version 2 is disabled.
VTP password	None.
VTP pruning	Disabled.

VTP Configuration Guidelines

These sections describe guidelines you should follow when implementing VTP in your network.

Domain Names

When configuring VTP for the first time, you must always assign a domain name. You must configure all switches in the VTP domain with the same domain name. Switches in VTP transparent mode do not exchange VTP messages with other switches, and you do not need to configure a VTP domain name for them.

Note If NVRAM and DRAM storage is sufficient, all switches in a VTP domain should be in VTP server mode.

Passwords

You can configure a password for the VTP domain, but it is not required. If you do configure a domain password, all domain switches must share the same password and you must configure the password on each switch in the management domain. Switches without a password or with the wrong password reject VTP advertisements.

If you configure a VTP password for a domain, a switch that is booted without a VTP configuration does not accept VTP advertisements until you configure it with the correct password. After the configuration, the switch accepts the next VTP advertisement that uses the same password and domain name in the advertisement.

If you are adding a new switch to an existing network with VTP capability, the new switch learns the domain name only after the applicable password has been configured on it.

Caution When you configure a VTP domain password, the management domain does not function properly if you do not assign a management domain password to each switch in the domain.

VTP Version

Follow these guidelines when deciding which VTP version to implement:

- All switches in a VTP domain must run the same VTP version.
- A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 if version 2 is disabled on the version 2-capable switch (version 2 is disabled by default).
- Do not enable VTP version 2 on a switch unless all of the switches in the same VTP domain are version-2-capable. When you enable version 2 on a switch, all of the version-2-capable switches in the domain enable version 2. If there is a version 1-only switch, it does not exchange VTP information with switches with version 2 enabled.
- If there are TrBRF and TrCRF Token Ring networks in your environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.
- Enabling or disabling VTP pruning on a VTP server enables or disables VTP pruning for the entire VTP domain.
- Making VLANs pruning-eligible or pruning-ineligible affects pruning eligibility for those VLANs on that device only (not on all switches in the VTP domain.)

Configuration Requirements

When you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements. For more information, see the "Understanding VLAN Trunks" section on page 9-22.

You can configure VTP by entering commands in the VLAN configuration mode. When you enter the **exit** command in VLAN configuration mode, it applies all the commands that you entered. VTP messages are sent to other switches in the VTP domain, and the privileged EXEC mode prompt appears.

Note The Cisco IOS **end** and **Ctrl-Z** commands are not supported in VLAN configuration mode.

For more configuration guidelines, see the "VLAN Configuration Guidelines" section on page 9-16.

Configuring a VTP Server

When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.

Beginning in privileged EXEC mode, follow these steps to configure the switch as a VTP server:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp server**	Configure the switch for VTP server mode (the default).
Step 3	**vtp domain** *domain-name*	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	**vtp password** *password-value*	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
Step 5	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 6	**show vtp status**	Verify your entries in the *VTP Operating Mode* and the *VTP Domain Name* fields of the display.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

To return the switch to a no-password state, use the **no vtp password** VLAN configuration command.

This example shows how to configure the switch as a VTP server with the domain name *eng_group* and verify the configuration:

```
Switch# vlan database
Switch(vlan)# vtp domain eng_group
Switch(vlan)# exit
APPLY completed.
Exiting....
Switch# show vtp status
VTP Version                     : 2
Configuration Revision          : 211
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Server
VTP Domain Name                 : eng_group
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x31 0xB3 0xCD 0xEF 0x34 0xD2 0x44 0xAD
Configuration last modified by 172.20.135.204 at 3-1-93 00:05:51
Local updater ID is 172.20.135.202 on interface Vl1 (lowest numbered VLAN interface found)
```

Configuring a VTP Client

When a switch is in VTP client mode, you cannot change its VLAN configuration. The client switch receives VTP updates from a VTP server in the VTP domain and then modifies its configuration accordingly.

Beginning in privileged EXEC mode, follow these steps to configure the switch for VTP client mode:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp client**	Configure the switch for VTP client mode. The default setting is VTP server.
Step 3	**vtp domain** *domain-name*	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. This should be the same domain name as the VTP server. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	**vtp password** *password-value*	(Optional) Assign a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
Step 5	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 6	**show vtp status**	Verify your entries in the *VTP Operating Mode* field of the display.

Use the **no vtp client** VLAN configuration command to return the switch to VTP server mode. To return the switch to a no-password state, use the **no vtp password** VLAN configuration command. When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

Disabling VTP (VTP Transparent Mode)

When you configure the switch for VTP transparent mode, you disable VTP on the switch. The switch then does not send VTP updates and does not act on VTP updates received from other switches. However, a VTP transparent switch running VTP version 2 does forward received VTP advertisements on all of its trunk links.

Beginning in privileged EXEC mode, follow these steps to configure the switch for VTP transparent mode:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp transparent**	Configure the switch for VTP transparent mode. The default setting is VTP server. This step disables VTP on the switch.

	Command	Purpose
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vtp status**	Verify your entries in the *VTP Operating Mode* field of the display.

To return the switch to VTP server mode, use the **no vtp transparent** VLAN configuration command.

Enabling VTP Version 2

VTP version 2 is disabled by default on VTP version 2-capable switches. When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2.

Caution VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Every switch in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.

Note In TrCRF and TrBRF Token ring environments, you must enable VTP version 2 for Token Ring VLAN switching to function properly. For Token Ring and Token Ring-Net media, VTP version 2 must be disabled.

For more information on VTP version configuration guidelines, see the "VTP Version" section on page 9-9.

Beginning in privileged EXEC mode, follow these steps to enable VTP version 2:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp v2-mode**	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vtp status**	Verify that VTP version 2 is enabled in the *VTP V2 Mode* field of the display.

To disable VTP version 2, use the **no vtp v2-mode** VLAN configuration command**.**

Enabling VTP Pruning

Pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the destination devices. You enable VTP pruning on a switch in VTP server mode.

Beginning in privileged EXEC mode, follow these steps to enable VTP pruning in the management domain:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp pruning**	Enable pruning in the VTP administrative domain. By default, pruning is disabled. You need to enable pruning on only one switch in VTP server mode.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vtp status**	Verify your entries in the *VTP Pruning Mode* field of the display.

Pruning is supported with VTP version 1 and version 2. If you enable pruning on the VTP server, it is enabled for the entire VTP domain.

Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on trunk ports. To change the pruning-eligible VLANs, see the "Changing the Pruning-Eligible List" section on page 9-28.

To disable VTP pruning, use the **no vtp pruning** VLAN configuration command.

Monitoring VTP

You monitor VTP by displaying VTP configuration information: the domain name, the current VTP revision, and the number of VLANs. You can also display statistics about the advertisements sent and received by the switch.

Beginning in privileged EXEC mode, follow these steps to monitor VTP activity:

Table 9-4 VTP Monitoring Commands

Command	Purpose
show vtp status	Display the VTP switch configuration information.
show vtp counters	Display counters about VTP messages that have been sent and received.

This is an example of output from the **show vtp status** privileged EXEC command:

```
Switch# show vtp status
VTP Version                     : 2
Configuration Revision          : 5
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 69
VTP Operating Mode              : Server
VTP Domain Name                 : test
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x59 0xBA 0x92 0xA4 0x74 0xD5 0x42 0x29
Configuration last modified by 0.0.0.0 at 3-1-93 00:18:42
Local updater ID is 10.1.1.59 on interface Vl1 (lowest numbered VLAN interface found)
```

This is an example of output from the **show vtp counters** privileged EXEC command:

```
Switch# show vtp counters
VTP statistics:
Summary advertisements received    : 0
Subset advertisements received     : 0
Request advertisements received    : 0
Summary advertisements transmitted : 0
Subset advertisements transmitted  : 0
Request advertisements transmitted : 0
Number of config revision errors   : 0
Number of config digest errors     : 0
Number of V1 summary errors        : 0


VTP pruning statistics:

Trunk            Join Transmitted Join Received    Summary advts received from
                                                   non-pruning-capable device
---------------- ---------------- ---------------- --------------------------
```

VLANs in the VTP Database

You can set these parameters when you create a new VLAN or modify an existing VLAN in the VTP database:

- VLAN ID
- VLAN name
- VLAN type (Ethernet, Fiber Distributed Data Interface [FDDI], FDDI network entity title [NET], TrBRF, or TrCRF, Token Ring, Token Ring-Net)
- VLAN state (active or suspended)
- Maximum transmission unit (MTU) for the VLAN
- Security Association Identifier (SAID)
- Bridge identification number for TrBRF VLANs
- Ring number for FDDI and TrCRF VLANs
- Parent VLAN number for TrCRF VLANs
- Spanning Tree Protocol (STP) type for TrCRF VLANs
- VLAN number to use when translating from one VLAN type to another

The "Default VLAN Configuration" section on page 9-15 lists the default values and possible ranges for each VLAN media type.

Token Ring VLANs

Although the Catalyst 3550 switches do not support Token Ring connections, a remote device such as a Catalyst 5000 series switch with Token Ring connections could be managed from one of the supported switches. Switches running VTP version 2 advertise information about these Token Ring VLANs:

- Token Ring TrBRF VLANs
- Token Ring TrCRF VLANs

For more information on configuring Token Ring VLANs, see the *Catalyst 5000 Series Software Configuration Guide*.

Default VLAN Configuration

Table 9-5 shows the default configuration for Ethernet VLANs.

Note The switch supports Ethernet interfaces exclusively. Because FDDI and Token Ring VLANs are not locally supported, you only configure FDDI and Token Ring media-specific characteristics for VTP global advertisements to other switches.

Table 9-5 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1–1005
VLAN name	default	No range
802.10 SAID	101001	1–4294967294
MTU size	1500	1500–18190
Translational bridge 1	1002	0–1005
Translational bridge 2	1003	0–1005
VLAN state	active	active, suspend

VLAN Configuration Guidelines

Follow these guidelines when creating and modifying VLANs in your network:

- The Catalyst 3550 switch supports 1005 VLANs in VTP client, server, and transparent modes. VLANs are identified with a number between 1 and 1001. VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs.
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. If the switch is a VTP server, you must define a VTP domain.
- Catalyst 3550 switches do not support Token Ring or FDDI media. The switch does not forward FDDI, FDDI-Net, TrCRF, or TrBRF traffic, but it does propagate the VLAN configuration through VTP.
- The switch supports 128 spanning-tree instances. If a switch has more active VLANs than supported spanning-tree instances, STP can be enabled on 128 VLANs and is disabled on the remaining VLANs. If you have already used all available spanning-tree instances on a switch, adding another VLAN anywhere in the VTP domain creates a VLAN on that switch that is not running spanning-tree. If you have the default allowed list on the trunk ports of that switch (which is to allow all VLANs), the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that would not be broken, particularly if there are several adjacent switches that all have run out of spanning-tree instances. You can prevent this possibility by setting allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances.
- If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.
- A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, the port mode is not changed.

Configuring VLANs in the VTP Database

You can add, modify or remove VLAN configurations in the VTP database by using the CLI VLAN configuration mode. VTP globally propagates these VLAN changes throughout the VTP domain.

In VTP server or transparent mode, commands to add, change, and delete VLANs are written to the file vlan.dat, and you can display them by entering the **show vlan** privileged EXEC command. The vlan.dat file is stored in NVRAM.

Caution You can cause inconsistency in the VLAN database if you attempt to manually delete the vlan.dat file. If you want to modify the VLAN configuration or VTP, use the **vlan database** privileged EXEC command to enter VLAN configuration mode as described in the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

You use the interface configuration mode to define the port membership mode and to add and remove ports from VLANs. The results of these commands are written to the running-configuration file, and you can display the file by entering the **show running-config** privileged EXEC command.

Note VLANs can be configured to support a number of parameters that are not discussed in detail in this section. For complete information on the commands and parameters that control VLAN configuration, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Adding an Ethernet VLAN

Each Ethernet VLAN has a unique, 4-digit ID that can be a number from 1 to 1001. To add a VLAN to the VLAN database, assign a number and name to the VLAN. For the list of default parameters that are assigned when you add a VLAN, see the "Default VLAN Configuration" section on page 9-15.

Beginning in privileged EXEC mode, follow these steps to add an Ethernet VLAN:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vlan** *vlan-id* **name** *vlan-name*	Add an Ethernet VLAN by assigning a number to it. If no name is entered for the VLAN, the default is to append the *vlan-id* to the word VLAN. For example, VLAN0004 could be a default VLAN name for VLAN 4.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan name** *vlan-name* or **show vlan**	Verify your entries.

To return the VLAN name to the default setting, use the **no vlan** *vlan-id* **name** VLAN configuration command.

This example shows how to add Ethernet VLAN 20 to the VLAN database and name it *test20*:

```
Switch# vlan database
Switch(vlan)# vlan 20 name test20
Switch(vlan)# exit
APPLY completed.
Exiting....
Switch# show vlan name test20

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
20   test20                           active

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
20   enet  100020     1500  -      -      -        -    -        0      0
```

Modifying an Ethernet VLAN

Beginning in privileged EXEC mode, follow these steps to modify an Ethernet VLAN:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vlan** *vlan-id* **mtu** *mtu-size*	Identify the VLAN, and change the MTU size.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan id** *vlan-id* or **show vlan**	Verify your entries.

To return the VLAN to the default MTU setting, use the **no vlan** *vlan-id* **mtu** VLAN configuration command.

This example shows how to verify a VLAN configuration:

```
Switch# show vlan id 20

VLAN Name                             Status    Ports
show vlan---- -------------------------------- --------- -------------------------------
20   VLAN0020                         active

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
20   enet  100020     1500  -      -      -        -    -        0      0
```

Deleting a VLAN from the Database

When you delete a VLAN from a switch that is in VTP server mode, the VLAN is removed from all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch.

You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.

Caution When you delete a VLAN, any ports assigned to that VLAN become inactive. They remain associated with the VLAN (and thus inactive) until you assign them to a new VLAN.

Beginning in privileged EXEC mode, follow these steps to delete a VLAN on the switch:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**no vlan** *vlan-id*	Remove the VLAN by entering the VLAN ID.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan brief**	Verify the VLAN removal.

This example shows how to verify a VLAN removal by using the **show vlan brief** privileged EXEC command:

```
Switch# show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Gi0/1, Gi0/2, Gi0/3, Gi0/4
                                                Gi0/5, Gi0/7, Gi0/8, Gi0/9
                                                Gi0/10, Gi0/11, Gi0/12
11   VLAN0011                         active
20   VLAN0020                         active
129  VLAN0129                         active
1002 fddi-default                     active
1003 token-ring-default               active
1004 fddinet-default                  active
1005 trnet-default                    active
```

Assigning Static-Access Ports to a VLAN

You can assign a static-access port to a VLAN without having VTP globally propagate VLAN configuration information (VTP is disabled).

Note If you assign an interface to a VLAN that does not exist, the new VLAN is created. (See the "Adding an Ethernet VLAN" section on page 9-17.)

Beginning in privileged EXEC mode, follow these steps to assign a port to a VLAN in the VTP database:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**interface** *interface-id*	Enter the interface to be added to the VLAN.
Step 3	**switchport mode access**	Define the VLAN membership mode for the port (Layer 2 access port).
Step 4	**switchport access vlan** *vlan-id*	Assign the port to a VLAN.
Step 5	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 6	**show running-config interface** *interface-id*	Verify the VLAN membership mode of the interface.
Step 7	**show interfaces** *interface-id* **switchport**	Verify your entries in the *Administrative Mode* and the *Access Mode VLAN* fields of the display.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/1 as an access port in VLAN 2:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
Switch#
```

These examples show how to verify the configuration:

```
Switch# show running-config interface gigabitethernet0/1
Building configuration...

Current configuration : 74 bytes
!
interface GigabitEthernet0/1
 no ip address
 snmp trap link-status
end

Switch# show interfaces gigabitethernet0/1 switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false
Unknown unicast blocked: false
Unknown multicast blocked: false

Broadcast Suppression Level: 100
Multicast Suppression Level: 100
Unicast Suppression Level: 100
```

Displaying VLANs in the VTP Database

Use the **show vlan** privileged EXEC command to display a list of VLANs in the database, including status, ports, and configuration:

```
Switch# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Gi0/1, Gi0/2, Gi0/3, Gi0/4
                                                Gi0/7, Gi0/8, Gi0/9, Gi0/11
                                                Gi0/12
20   VLAN0020                         active
21   VLAN0021                         active
22   VLAN0022                         active
27   VLAN0027                         active
31   VLAN0031                         active
1002 fddi-default                     active
1003 trcrf-default                    active
1004 fddinet-default                  active
1005 trbrf-default                    active

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        1002   1003
20   enet  100020     1500  -      -      -        -    -        0      0
21   enet  100021     1500  -      -      -        -    -        0      0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
22   enet  100022     1500  -      -      -        -    -        0      0
27   enet  100027     1500  -      -      -        -    -        0      0
31   enet  100031     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        1      1003
1003 trcrf 101003     4472  1005   3276   -        -    srb      1      1002
1004 fdnet 101004     1500  -      -      1        ibm  -        0      0
1005 trbrf 101005     4472  -      -      15       ibm  -        0      0

VLAN AREHops STEHops Backup CRF
---- ------- ------- ----------
1003 7       7       off
```

Use the **show vlan brief** privileged EXEC command to display a list of VLANs in the database with status and port information but without configuration information:

```
Switch# show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Gi0/1, Gi0/2, Gi0/3, Gi0/4
                                                Gi0/7, Gi0/8, Gi0/9, Gi0/11
                                                Gi0/12
20   VLAN0020                         active
21   VLAN0021                         active
22   VLAN0022                         active
27   VLAN0027                         active
31   VLAN0031                         active
1002 fddi-default                     active
1003 trcrf-default                    active
1004 fddinet-default                  active
1005 trbrf-default                    active
```

Understanding VLAN Trunks

These sections describe how VLAN trunks function on the switch:

- Trunking Overview, page 9-22
- Encapsulation Types, page 9-23
- Default Layer 2 Ethernet Interface VLAN Configuration, page 9-24

Trunking Overview

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Trunks carry the traffic of multiple VLANs over a single link, and you can extend VLANs across an entire network. The 100BASE-T and Gigabit Ethernet trunks carry traffic for multiple VLANs over a single link.

Two trunking encapsulations are available on all Ethernet interfaces:

- Inter-Switch Link (ISL)—ISL is Cisco-proprietary trunking encapsulation.
- 802.1Q—802.1Q is industry-standard trunking encapsulation.

Figure 9-4 shows a network of switches that are connected by ISL trunks.

Figure 9-4 Switches in an ISL Trunking Environment

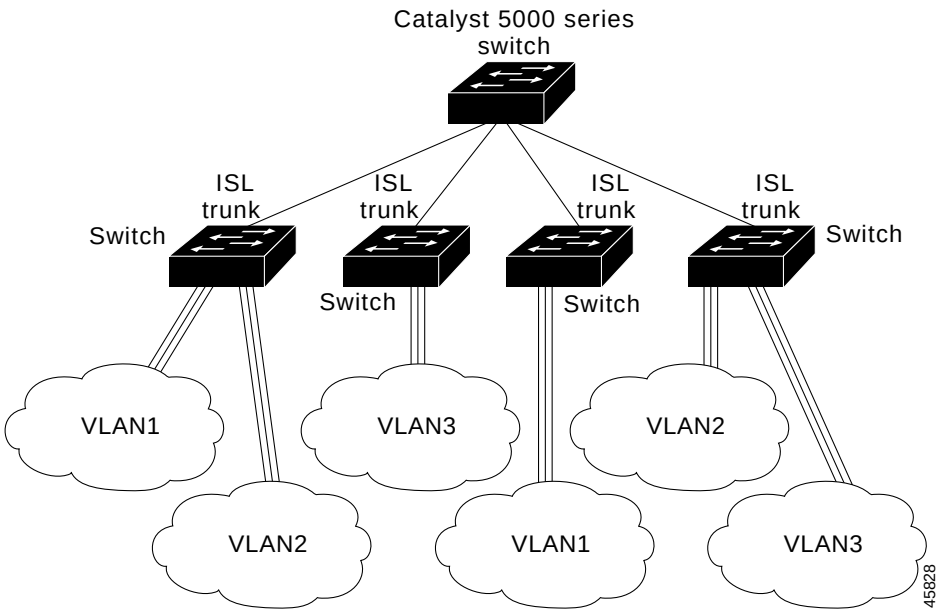

You can configure a trunk on a single Ethernet interface or on an EtherChannel bundle. For more information about EtherChannel, see Chapter 21, "Configuring EtherChannel."

Ethernet trunk interfaces support different trunking modes (see Table 9-6). You can specify whether the trunk uses ISL or 802.1Q encapsulation or if the encapsulation type is autonegotiated. To autonegotiate trunking, the interfaces must be in the same VTP domain. Use the **trunk** or **nonegotiate** keywords to force interfaces in different domains to trunk. Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which supports autonegotiation of both ISL and 802.1Q trunks.

Note DTP is a point-to-point protocol. However, some internetworking devices might forward DTP frames improperly. To avoid this, ensure that interfaces connected to devices that do not support DTP are configured with the **access** keyword if you do not intend to trunk across those links. To enable trunking to a device that does not support DTP, use the **nonegotiate** keyword to cause the interface to become a trunk but to not generate DTP frames.

Table 9-6 Layer 2 Interface Modes

Mode	Function
switchport mode access	Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface even if the neighboring interface is not a trunk interface.
switchport mode dynamic desirable	Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to *trunk*, *desirable*, or *auto* mode. The default switch-port mode for all Ethernet interfaces is **dynamic desirable**.
switchport mode dynamic auto	Makes the interface able to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to *trunk* or *desirable* mode.
switchport mode trunk	Puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
switchport nonegotiate	Puts the interface into permanent trunking mode but prevents the interface from generating DTP frames. You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

Encapsulation Types

Table 9-7 lists the Ethernet trunk encapsulation types and keywords.

Table 9-7 Ethernet Trunk Encapsulation Types

Encapsulation	Function
switchport trunk encapsulation isl	Specifies ISL encapsulation on the trunk link.
switchport trunk encapsulation dot1q	Specifies 802.1Q encapsulation on the trunk link.
switchport trunk encapsulation negotiate	Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface.

Note The switch does not support Layer 3 trunks; you cannot configure subinterfaces or use the **encapsulation** keyword on Layer 3 interfaces. The switch does support Layer 2 trunks and Layer 3 VLAN interfaces, which provide equivalent capabilities.

The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected interfaces determine whether a link becomes an ISL or 802.1Q trunk.

802.1Q Configuration Considerations

802.1Q trunks impose these limitations on the trunking strategy for a network:

- In a network of Cisco switches connected through 802.1Q trunks, the switches maintain one instance of spanning tree for each VLAN allowed on the trunks. Non-Cisco devices might support one spanning-tree instance for all VLANs.

  When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch combines the spanning-tree instance of the VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch. However, spanning-tree information for each VLAN is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.
- Make sure the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning-tree loops might result.
- Disabling STP on the native VLAN of an 802.1Q trunk without disabling STP on every VLAN in the network can potentially cause spanning-tree loops. We recommend that you leave STP enabled on the native VLAN of an 802.1Q trunk or disable STP on every VLAN in the network. Make sure your network is loop-free before disabling STP.

Default Layer 2 Ethernet Interface VLAN Configuration

Table 9-8 shows the default Layer 2 Ethernet interface VLAN configuration.

Table 9-8 Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Interface mode	**switchport mode dynamic desirable**
Trunk encapsulation	**switchport trunk encapsulation negotiate**
Allowed VLAN range	VLANs 1–1005
VLAN range eligible for pruning	VLANs 2–1001
Default VLAN (for access ports)	VLAN 1
Native VLAN (for 802.1Q trunks)	VLAN 1

Configuring an Ethernet Interface as a Trunk Port

Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

This section includes these procedures for configuring an Ethernet interface as a trunk port on the switch:

- Configuring a Trunk Port, page 9-25
- Defining the Allowed VLANs on a Trunk, page 9-27
- Changing the Pruning-Eligible List, page 9-28
- Configuring the Native VLAN for Untagged Traffic, page 9-29

Note By default, an interface is in Layer 2 mode. The default mode for Layer 2 interfaces is **switchport mode dynamic desirable**. If the neighboring interface supports trunking and is configured to allow trunking, the link is a Layer 2 trunk or, if the interface is in Layer 3 mode, it becomes a Layer 2 trunk when you enter the **switchport** interface configuration command. By default, trunks negotiate encapsulation. If the neighboring interface supports ISL and 802.1Q encapsulation and both interfaces are set to negotiate the encapsulation type, the trunk uses ISL encapsulation.

Configuring a Trunk Port

Beginning in privileged EXEC mode, follow these steps to configure a port as an ISL or 802.1Q trunk port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter the interface configuration mode and the port to be configured for trunking.
Step 3	**switchport trunk encapsulation** {**isl** | **dot1q** | **negotiate**}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate (the default) with the neighboring interface for encapsulation type. You must configure each end of the link with the same encapsulation type.
Step 4	**switchport mode** {**dynamic** {**auto** | **desirable**} | **trunk**}	Configure the interface as a Layer 2 trunk (required only if the interface is a Layer 2 access port or to specify the trunking mode). • **dynamic auto**—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. • **dynamic desirable**—Set the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode. • **trunk**—Set the interface in permanent trunking mode and negotiate to convert the link to a trunk link even if the neighboring interface is not a trunk interface.
Step 5	**switchport access vlan** *vlan-id*	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
Step 6	**switchport trunk native vlan** *vlan-id*	For 802.1Q trunks, specify the native VLAN.
Step 7	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 8	**show interfaces** *interface-id* **switchport**	Display the switchport configuration of the interface in the *Administrative Mode* and the *Administrative Trunking Encapsulation* fields of the display.
Step 9	**show interfaces** *interface-id* **trunk**	Display the trunk configuration of the interface.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To reset all trunking characteristics of a trunking interface to the defaults, use the **no switchport trunk** interface configuration command. Use the **no switchport mode** and **no switchport access** interface configuration commands to return to the default settings.

This example shows how to configure the Gigabit Ethernet interface 0/4 as an 802.1Q trunk. This example assumes that the neighbor interface is configured to support 802.1Q trunking:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# switchport mode dynamic desirable
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# end
Switch#
```

These examples show how to verify the configuration:

```
Switch# show running-config interface gigabitethernet0/4
Building configuration...

Current configuration : 112 bytes
!
interface GigabitEthernet0/4
 switchport trunk encapsulation dot1q
 no ip address
 snmp trap link-status
end

Switch# show interfaces gigabitethernet0/4 switchport
Name: Gi0/4
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false
Unknown unicast blocked: false
Unknown multicast blocked: false

Broadcast Suppression Level: 100
Multicast Suppression Level: 100
Unicast Suppression Level: 100
```

In this example, the encapsulation method is ISL:

```
Switch# show interfaces gigabitethernet0/4 trunk

Port      Mode         Encapsulation  Status        Native vlan
Gi0/4     desirable    n-isl          trunking      1

Port      Vlans allowed on trunk
Gi0/4     1-1005

Port      Vlans allowed and active in management domain
Gi0/4     1,10-1000

Port      Vlans in spanning tree forwarding state and not pruned
Gi0/4     1,10-1000
```

Defining the Allowed VLANs on a Trunk

By default, a trunk port sends traffic to and receives traffic from all VLANs in the VLAN database. All VLANs, VLANs 1 to 1005, are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **switchport trunk allowed vlan remove** *vlan-list* interface configuration to remove specific VLANs from the allowed list.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Beginning in privileged EXEC mode, follow these steps to modify the allowed list of an ISL or 802.1Q trunk:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode and the port to be added to the VLAN.
Step 3	**switchport trunk allowed vlan** {**add** | **except** | **none** | **remove**} *vlan-list*	(Optional) Configure the list of VLANs allowed on the trunk. For explanations about using the **add**, **except**, **none**, and **remove** keywords, refer to *Catalyst 3550 Multilayer Switch Command Reference* for this release. The *vlan-list* parameter is either a single VLAN number from 1 to 1005 or a range of VLANs described by two VLAN numbers, the lesser one first, separated by a hyphen. Do not enter any spaces between comma-separated VLAN parameters or in hyphen-specified ranges. All VLANs are allowed by default. You cannot remove any of the default VLANs from a trunk.
Step 4	**end**	Return to privileged EXEC.
Step 5	**show interfaces** *interface-id* **switchport**	Verify your entries in the *Trunking VLANs Enabled* field of the display.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default allowed VLAN list of all VLANs, use the **no switchport trunk allowed vlan** interface configuration command.

This example shows how to remove VLAN 2 from the allowed VLAN list and verify the configuration.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport trunk allowed vlan remove 2
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/1 switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 1,3-1005
Pruning VLANs Enabled: 2-1001

Protected: false
Unknown unicast blocked: false
Unknown multicast blocked: false

Broadcast Suppression Level: 100
Multicast Suppression Level: 100
Unicast Suppression Level: 100
```

Changing the Pruning-Eligible List

The pruning-eligible list applies only to trunk ports. Each trunk port has its own eligibility list. VTP pruning must be enabled for this procedure to take effect. The "Enabling VTP Pruning" section on page 9-13 describes how to enable VTP pruning.

Beginning in privileged EXEC mode, follow these steps to remove VLANs from the pruning-eligible list on a trunk port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and select the trunk port for which VLANs should be pruned.
Step 3	**switchport trunk pruning vlan** {**add** | **except** | **none** | **remove**} *vlan-list* [,*vlan*[,*vlan*[,,,]]	Configure the list of VLANs allowed to be pruned from the trunk. (See the "VTP Pruning" section on page 9-6). For explanations about using the **add**, **except**, **none**, and **remove** keywords, refer to *Catalyst 3550 Multilayer Switch Command Reference* for this release. Separate nonconsecutive VLAN IDs with a comma and no spaces; use a hyphen to designate a range of IDs. Valid IDs are from 2 to 1001. VLANs that are pruning-ineligible receive flooded traffic. The default list of VLANs allowed to be pruned contains all VLANs.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show interfaces** *interface-id* **switchport**	Verify your entries in the *Pruning VLANs Enabled* field of the display.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default pruning-eligible list of all VLANs, use the **no switchport trunk pruning vlan** interface configuration command.

Configuring the Native VLAN for Untagged Traffic

A trunk port configured with 802.1Q tagging can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN configured for the port. The native VLAN is VLAN 1 by default.

Note The native VLAN can be assigned any VLAN ID.

For information about 802.1Q configuration issues, see the "Encapsulation Types" section on page 9-23.

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on an 802.1Q trunk:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and define the interface that is configured as the 802.1Q trunk.
Step 3	**switchport trunk native vlan** *vlan-id*	Configure the VLAN that is sending and receiving untagged traffic on the trunk port. Valid IDs are from 1 to 1001.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show interfaces** *interface-id* **switchport**	Verify your entries in the *Trunking Native Mode VLAN* field.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default native VLAN, VLAN 1, use the **no switchport trunk native vlan** interface configuration command.

If a packet has a VLAN ID that is the same as the outgoing port native VLAN ID, the packet is sent untagged; otherwise, the switch sends the packet with a tag.

Load Sharing Using STP

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches. For more information, see Chapter 10, "Configuring STP."

Load Sharing Using STP Port Priorities

When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

Figure 9-5 shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 10 on Trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on Trunk 1.
- VLANs 3 through 6 are assigned a port priority of 10 on Trunk 2.
- VLANs 8 through 10 retain the default port priority of 128 on Trunk 2.

In this way, Trunk 1 carries traffic for VLANs 8 through 10, and Trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

Figure 9-5 Load Sharing by Using STP Port Priorities

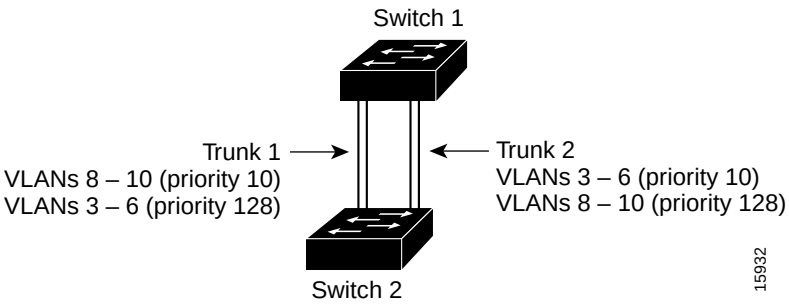

Configuring STP Port Priorities and Load Sharing

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 9-5:

	Command	Purpose
Step 1	**vlan database**	On Switch 1, enter VLAN configuration mode.
Step 2	**vtp domain** *domain-name*	Configure a VTP administrative domain. The domain name can be from 1 to 32 characters.
Step 3	**vtp server**	Configure Switch 1 as the VTP server.
Step 4	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 5	**show vtp status**	Verify the VTP configuration on both Switch 1 and Switch 2. In the display, check the *VTP Operating Mode* and the *VTP Domain Name* fields.

	Command	Purpose
Step 6	**show vlan**	Verify that the VLANs exist in the database on Switch 1.
Step 7	**configure terminal**	Enter global configuration mode.
Step 8	**interface gigabitethernet 0/1**	Enter interface configuration mode, and define Gigabit Ethernet port 0/1 as the interface to be configured as a trunk.
Step 9	**switchport trunk encapsulation** {**isl** | **dot1q** | **negotiate**}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate with the neighboring interface. You must configure each end of the link with the same encapsulation type.
Step 10	**switchport mode trunk**	Configure the port as a trunk port.
Step 11	**end**	Return to privilege EXEC mode.
Step 12	**show interfaces gigabitethernet0/1 switchport**	Verify the VLAN configuration.
Step 13		Repeat Steps 7 through 11 on Switch 1 for interface Gi0/2.
Step 14		Repeat Steps 7 through 11 on Switch 2 to configure the trunk ports on interface Gi0/1 and Gi0/2.
Step 15	**show vlan**	When the trunk links come up, VTP passes the VTP and VLAN information to Switch 2. Verify that Switch 2 has learned the VLAN configuration.
Step 16	**configure terminal**	Enter global configuration mode on Switch 1.
Step 17	**interface gigabitethernet0/1**	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 18	**spanning-tree vlan 8 port-priority 10**	Assign the port priority of 10 for VLAN 8.
Step 19	**spanning-tree vlan 9 port-priority 10**	Assign the port priority of 10 for VLAN 9.
Step 20	**spanning-tree vlan 10 port-priority 10**	Assign the port priority of 10 for VLAN 10.
Step 21	**exit**	Return to global configuration mode.
Step 22	**interface gigabitethernet0/2**	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 23	**spanning-tree vlan 3 port-priority 10**	Assign the port priority of 10 for VLAN 3.
Step 24	**spanning-tree vlan 4 port-priority 10**	Assign the port priority of 10 for VLAN 4.
Step 25	**spanning-tree vlan 5 port-priority 10**	Assign the port priority of 10 for VLAN 5.
Step 26	**spanning-tree vlan 6 port-priority 10**	Assign the port priority of 10 for VLAN 6.
Step 27	**end**	Return to privileged EXEC mode.
Step 28	**show running-config**	Verify your entries.
Step 29	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Load Sharing Using STP Path Cost

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs. The VLANs keep the traffic separate. Because no loops exist, STP does not disable the ports, and redundancy is maintained in the event of a lost link.

In Figure 9-6, Trunk ports 1 and 2 are 100BASE-T ports. The path costs for the VLANs are assigned as follows:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 100BASE-T path cost on Trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 100BASE-T path cost on Trunk port 2 of 19.

Figure 9-6 Load-Sharing Trunks with Traffic Distributed by Path Cost

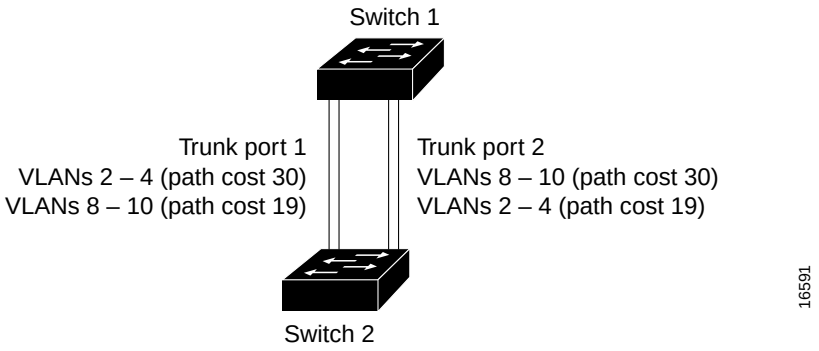
Switch 1
Trunk port 1
VLANs 2 – 4 (path cost 30)
VLANs 8 – 10 (path cost 19)
Trunk port 2
VLANs 8 – 10 (path cost 30)
VLANs 2 – 4 (path cost 19)
Switch 2
16591

Configuring STP Path Costs and Load Sharing

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 9-6:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode on Switch 1.
Step 2	**interface fastethernet 0/1**	Enter interface configuration mode, and define Fast Ethernet port 0/1 as the interface to be configured as a trunk.
Step 3	**switchport trunk encapsulation {isl | dot1q | negotiate}**	Configure the port to support ISL or 802.1Q encapsulation. You must configure each end of the link with the same encapsulation type.
Step 4	**switchport mode trunk**	Configure the port as a trunk port. The trunk defaults to ISL trunking.
Step 5	**exit**	Return to global configuration mode.
Step 6		Repeat Steps 2 through 4 on Switch 1 interface Fast Ethernet 0/2.
Step 7	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 8	**show running-config**	Verify your entries. In the display, make sure that interfaces Fast Ethernet 0/1 and Fast Ethernet 0/2 are configured as trunk ports.
Step 9	**show vlan**	When the trunk links come up, Switch 1 receives the VTP information from the other switches. Verify that Switch 1 has learned the VLAN configuration.
Step 10	**configure terminal**	Enter global configuration mode.
Step 11	**interface fastethernet 0/1**	Enter interface configuration mode, and define Fast Ethernet port 0/1 as the interface to set the STP cost.
Step 12	**spanning-tree vlan 2 cost 30**	Set the spanning-tree path cost to 30 for VLAN 2.
Step 13	**spanning-tree vlan 3 cost 30**	Set the spanning-tree path cost to 30 for VLAN 3.
Step 14	**spanning-tree vlan 4 cost 30**	Set the spanning-tree path cost to 30 for VLAN 4.
Step 15	**end**	Return to global configuration mode.
Step 16		Repeat Steps 9 through 11 on Switch 1 interface Fast Ethernet 0/2, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.
Step 17	**exit**	Return to privileged EXEC mode.
Step 18	**show running-config**	Verify your entries. In the display, verify that the path costs are set correctly for interfaces Fast Ethernet 0/1 and 0/2.
Step 19	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Understanding VMPS

The Catalyst 3550 switch acts as a client to the VMPS and communicates with it through the VLAN Query Protocol (VQP). When the VMPS receives a VQP request from a client switch, it searches its database for a MAC-address-to-VLAN mapping. The server response is based on this mapping and whether or not the server is in secure mode. Secure mode determines whether the server shuts down the port when a VLAN is not allowed on it or just denies the port access to the VLAN.

In response to a request, the VMPS takes one of these actions:

- If the assigned VLAN is restricted to a group of ports, the VMPS verifies the requesting port against this group and responds as follows:
  - If the VLAN is allowed on the port, the VMPS sends the VLAN name to the client in response.
  - If the VLAN is not allowed on the port and the VMPS is not in secure mode, the VMPS sends an *access-denied* response.
  - If the VLAN is not allowed on the port and the VMPS is in secure mode, the VMPS sends a *port-shutdown* response.
- If the VLAN in the database does not match the current VLAN on the port and active hosts exist on the port, the VMPS sends an *access-denied* or a *port-shutdown* response, depending on the secure mode of the VMPS.

If the switch receives an *access-denied* response from the VMPS, it continues to block traffic from the MAC address to or from the port. The switch continues to monitor the packets directed to the port and sends a query to the VMPS when it identifies a new address. If the switch receives a *port-shutdown* response from the VMPS, it disables the port. The port must be manually re-enabled by using the CLI, CMS, or SNMP.

You can also use an explicit entry in the configuration table to deny access to specific MAC addresses for security reasons. If you enter the **none** keyword for the VLAN name, the VMPS sends an *access-denied* or *port-shutdown* response, depending on the VMPS secure mode setting.

Dynamic Port VLAN Membership

A dynamic (nontrunking) port on the switch can belong to only one VLAN. When the link comes up, the switch does not forward traffic to or from this port until the VMPS provides the VLAN assignment. The VMPS receives the source MAC address from the first packet of a new host connected to the dynamic port and attempts to match the MAC address to a VLAN in the VMPS database.

If there is a match, the VMPS sends the VLAN number for that port. If the client switch was not previously configured, it uses the domain name from the first VTP packet it receives on its trunk port from the VMPS. If the client switch was previously configured, it includes its domain name in the query packet to the VMPS to obtain its VLAN number. The VMPS verifies that the domain name in the packet matches its own domain name before accepting the request and responds to the client with the assigned VLAN number for the client. If there is no match, the VMPS either denies the request or shuts down the port (depending on the VMPS secure mode setting).

Multiple hosts (MAC addresses) can be active on a dynamic port if they are all in the same VLAN; however, the VMPS shuts down a dynamic port if more than 20 hosts are active on the port.

If the link goes down on a dynamic port, the port returns to an isolated state and does not belong to a VLAN. Any hosts that come online through the port are checked again through the VQP with the VMPS before the port is assigned to a VLAN.

VMPS Database Configuration File

The VMPS contains a database configuration file that you create. This ASCII text file is stored on a switch-accessible TFTP server that functions as a VMPS server. The file contains VMPS information, such as the domain name, the fallback VLAN name, and the MAC-address-to-VLAN mapping. The Catalyst 3550 switch cannot act as the VMPS, but you can use a Catalyst 5000 or Catalyst 6000 series switch as the VMPS.

You can configure a fallback VLAN name. If you connect a device with a MAC address that is not in the database, the VMPS sends the fallback VLAN name to the client. If you do not configure a fallback VLAN and the MAC address does not exist in the database, the VMPS sends an *access-denied* response. If the VMPS is in secure mode, it sends a *port-shutdown* response.

Whenever port names are used in the VMPS database configuration file, the server must use the switch convention for naming ports. For example, Gi0/4 is fixed Gigabit Ethernet port number 4. If the switch is a cluster member, the command switch adds the name of the switch before the type. For example, *es3%Gi0/4* refers to fixed Gigabit Ethernet port 4 on member switch 3. When port names are required, these naming conventions must be followed in the VMPS database configuration file when it is configured to support a cluster.

This example shows a example of a VMPS database configuration file as it appears on a Catalyst 6000 series switch. The file has these characteristics:

- The security mode is open.
- The default is used for the fallback VLAN.
- MAC address-to-VLAN name mappings—The MAC address of each host and the VLAN to which each host belongs is defined.
- Port groups are defined.
- VLAN groups are defined.
- VLAN port policies are defined for the ports associated with restricted VLANs.

```
!VMPS File Format, version 1.1
! Always begin the configuration file with
! the word “VMPS”
!
!vmps domain <domain-name>
! The VMPS domain must be defined.
!vmps mode {open | secure}
! The default mode is open.
!vmps fallback <vlan-name>
!vmps no-domain-req { allow | deny }
!
! The default value is allow.
vmps domain DSBU
vmps mode open
vmps fallback default
vmps no-domain-req deny
!
!
!MAC Addresses
!
vmps-mac-addrs
!
! address <addr> vlan-name <vlan_name>
!
address 0012.2233.4455 vlan-name hardware
address 0000.6509.a080 vlan-name hardware
address aabb.ccdd.eeff vlan-name Green
address 1223.5678.9abc vlan-name ExecStaff
address fedc.ba98.7654 vlan-name --NONE--
address fedc.ba23.1245 vlan-name Purple
!
!Port Groups
!
!vmps-port-group <group-name>
! device <device-id> { port <port-name> | all-ports }
!
vmps-port-group WiringCloset1
 device 198.92.30.32 port 0/2
 device 172.20.26.141 port 0/8
vmps-port-group “Executive Row”
 device 198.4.254.222 port 0/2
 device 198.4.254.222 port 0/3
 device 198.4.254.223 all-ports
!
!
!VLAN groups
!
!vmps-vlan-group <group-name>
! vlan-name <vlan-name>
```

```
!
vmps-vlan-group Engineering
vlan-name hardware
vlan-name software
!
!
!VLAN port Policies
!
!vmps-port-policies {vlan-name <vlan_name> | vlan-group <group-name> }
! { port-group <group-name> | device <device-id> port <port-name> }
!
vmps-port-policies vlan-group Engineering
 port-group WiringCloset1
vmps-port-policies vlan-name Green
 device 198.92.30.32 port 0/8
vmps-port-policies vlan-name Purple
 device 198.4.254.22 port 0/2
 port-group “Executive Row”
```

VMPS Configuration Guidelines

These guidelines and restrictions apply to dynamic port VLAN membership:

- You must configure the VMPS before you configure ports as dynamic.
- The communication between a cluster of switches and VMPS is managed by the command switch and includes port-naming conventions that are different from standard port names. For the cluster-based port-naming conventions, see the “VMPS Database Configuration File” section on page 9-34.
- When you configure a port as dynamic, the spanning-tree Port Fast feature is automatically enabled for that port. The Port Fast mode accelerates the process of bringing the port into the forwarding state. You can disable Port Fast mode on a dynamic port.
- If you try to enable 802.1X on a dynamic-access (VQP) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
- Trunk ports cannot be dynamic ports, but you can enter the **switchport access vlan dynamic** interface configuration command for a trunk port. In this case, the switch retains the setting and applies it if the port is later configured as an access port.

  You must turn off trunking on the port before the dynamic access setting takes effect.
- Dynamic ports cannot be monitor ports.
- Dynamic ports cannot be members of an EtherChannel group.
- A dynamic access port can participate in fallback bridging.
- The VTP management domain of the VMPS client and the VMPS server must be the same.

Default VMPS Configuration

Table 9-9 shows the default VMPS and dynamic port configuration on client switches.

Table 9-9 Default VMPS Client and Dynamic Port Configuration

Feature	Default Setting
VMPS domain server	None
VMPS reconfirm interval	60 minutes
VMPS server retry count	3
Dynamic ports	None configured

Configuring an Interface as a Layer 2 Dynamic Access Port

You configure dynamic VLANs by using the VMPS (server). The switch cannot be a VMPS server; the server can be a Catalyst 6000 switch or other similar device.

Entering the IP Address of the VMPS

You must first enter the IP address of the server to configure the switch as a client.

Note If the VMPS is being defined for a cluster of switches, enter the address on the command switch.

Beginning in privileged EXEC mode, follow these steps to enter the IP address of the VMPS:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**vmps server** *ipaddress* **primary**	Enter the IP address of the switch acting as the primary VMPS server.
Step 3	**vmps server** *ipaddress*	Enter the IP address of the switch acting as a secondary VMPS server. You can enter up to three secondary server addresses.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show vmps**	Verify your entries in the *VMPS Domain Server* field of the display.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

This is an example of output for the **show vmps** privileged EXEC command, used to verify the VMPS server IP address.

```
Switch# show vmps

VQP Client Status:
--------------------
VMPS VQP Version:   1
Reconfirm Interval: 60 min
Server Retry Count: 3
VMPS domain server: 172.20.128.86 (primary, current)
                    172.20.128.87

Reconfirmation status
---------------------
VMPS Action:         No Dynamic Port
```

Note The switch port that is connected to the VMPS server cannot be a dynamic access port. It can be either a static access port or a trunk port. See the "Configuring an Ethernet Interface as a Trunk Port" section on page 9-25.

Configuring Dynamic Access Ports on VMPS Clients

Caution Dynamic port VLAN membership is for end stations or hubs connected to end stations. Connecting dynamic access ports to other switches can cause a loss of connectivity.

Beginning in privileged EXEC mode, follow these steps to configure a dynamic access port on a VMPS client switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode and the switch port that is connected to the end station.
Step 3	**switchport mode access**	Set the port to access mode.
Step 4	**switchport access vlan dynamic**	Configure the port as eligible for dynamic VLAN membership. The dynamic access port must be connected to an end station.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show interfaces** *interface-id* **switchport**	Verify your entries in the *Operational Mode* field of the display.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To return an interface to its default switchport mode (dynamic desirable), use the **no switchport mode** interface configuration command. To reset the access mode to the default VLAN for the switch, use the **no switchport access** interface configuration command.

Reconfirming VLAN Memberships

Beginning in privileged EXEC mode, follow these steps to confirm the dynamic port VLAN membership assignments that the switch has received from the VMPS:

	Command	Purpose
Step 1	**vmps reconfirm**	Reconfirm dynamic port VLAN membership.
Step 2	**show vmps**	Verify the dynamic VLAN reconfirmation status.

Changing the Reconfirmation Interval

VMPS clients periodically reconfirm the VLAN membership information received from the VMPS. You can set the number of minutes after which reconfirmation occurs.

If you are configuring a member switch in a cluster, this parameter must be equal to or greater than the reconfirmation setting on the command switch.

Beginning in privileged EXEC mode, follow these steps to change the reconfirmation interval:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**vmps reconfirm** *minutes*	Enter the number of minutes between reconfirmations of the dynamic VLAN membership. Enter a number from 1 to 120. The default is 60 minutes.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show vmps**	Verify the dynamic VLAN reconfirmation status in the *Reconfirm Interval* field of the display.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps reconfirm** global configuration command.

Changing the Retry Count

Beginning in privileged EXEC mode, follow these steps to change the number of times that the switch attempts to contact the VMPS before querying the next server:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**vmps retry** *count*	Change the retry count. The retry range is from 1 to 10; the default is 3.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show vmps**	Verify your entry in the *Server Retry Count* field of the display.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps retry** global configuration command.

Administering and Monitoring the VMPS

You can display information about the VMPS by using the **show vmps** privileged EXEC command. The switch displays this information about the VMPS:

VMPS VQP Version	The version of VQP used to communicate with the VMPS. The switch queries the VMPS that is using VQP version 1.
Reconfirm Interval	The number of minutes the switch waits before reconfirming the VLAN-to-MAC-address assignments.
Server Retry Count	The number of times VQP resends a query to the VMPS. If no response is received after this many tries, the switch starts to query the secondary VMPS.
VMPS domain server	The IP address of the configured VLAN membership policy servers. The switch sends queries to the one marked *current*. The one marked *primary* is the primary server.
VMPS Action	The result of the most recent reconfirmation attempt. A reconfirmation attempt can occur automatically when the reconfirmation interval expired, or you can force it by entering the **vmps reconfirm** privileged EXEC command or its CMS or SNMP equivalent.

Troubleshooting Dynamic Port VLAN Membership

The VMPS shuts down a dynamic port under these conditions:

- The VMPS is in secure mode, and it does not allow the host to connect to the port. The VMPS shuts down the port to prevent the host from connecting to the network.
- More than 20 active hosts reside on a dynamic port.

To re-enable a disabled dynamic port, enter the **no shutdown** interface configuration command.

Dynamic Port VLAN Membership Configuration Example

Figure 9-7 shows a network with a VMPS server switch and VMPS client switches with dynamic ports. In this example, these assumptions apply:

- The VMPS server and the VMPS client are separate switches.
- The Catalyst 6000 series Switch 1 is the primary VMPS server.
- The Catalyst 6000 series Switch 3 and Switch 10 are secondary VMPS servers.
- End stations are connected to the Catalyst 3550 clients, Switch 2 and Switch 9.
- The database configuration file is stored on the TFTP server with the IP address 172.20.22.7.

Figure 9-7 Dynamic Port VLAN Membership Configuration

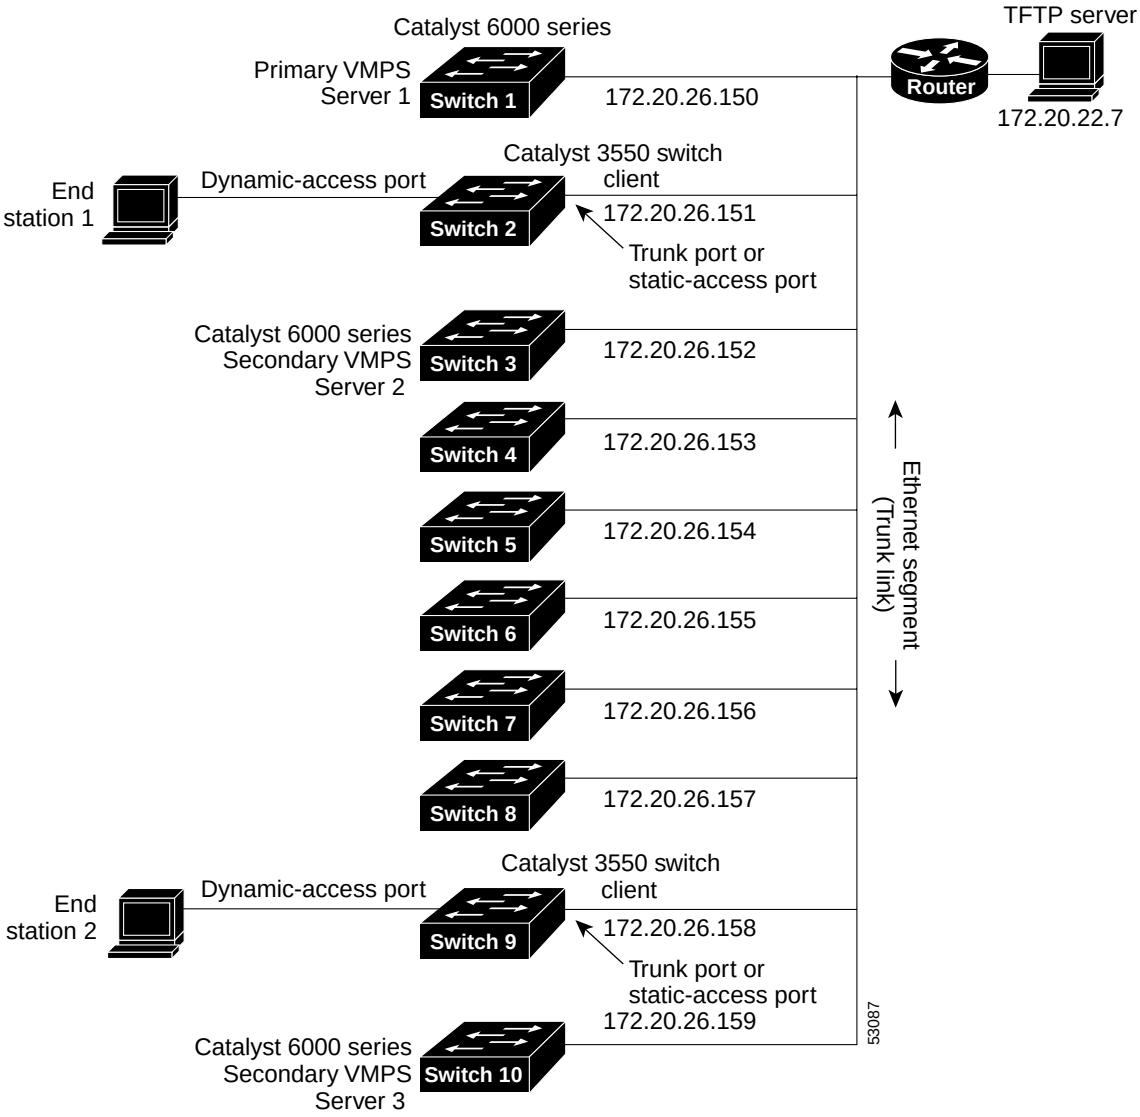

Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding Basic STP Features, page 10-1
- Understanding Advanced STP Features, page 10-10
- Configuring Basic STP Features, page 10-21
- Configuring Advanced STP Features, page 10-32

Understanding Basic STP Features

This section describes how basic STP features work. It includes this information:

- Supported STP Instances, page 10-2
- STP Overview, page 10-2
- Bridge ID, Switch Priority, and Extended System ID, page 10-3
- Election of the Root Switch, page 10-3
- Bridge Protocol Data Units, page 10-4
- STP Timers, page 10-5
- Creating the STP Topology, page 10-5
- STP Interface States, page 10-6
- STP Address Management, page 10-8
- STP and IEEE 802.1Q Trunks, page 10-8
- VLAN-Bridge STP, page 10-9
- STP and Redundant Connectivity, page 10-9
- Accelerated Aging to Retain Connectivity, page 10-10

For configuration information, see the "Configuring Basic STP Features" section on page 10-21.

For information about advanced STP features, see the "Understanding Advanced STP Features" section on page 10-10 and the "Configuring Advanced STP Features" section on page 10-32.

Supported STP Instances

This software release supports the per-VLAN spanning tree (PVST) and a maximum of 128 spanning-tree instances. If more VLANs are defined in the VLAN Trunking Protocol (VTP) than spanning-tree instances, you can enable STP on only 128 VLANs. The remaining VLANs operate with STP disabled.

If 128 instances of spanning tree are already in use, you can disable STP on one of the VLANs and then enable it on the VLAN where you want it to run. Use the **no spanning-tree vlan** *vlan-id* global configuration command to disable STP on a specific VLAN, and use the **spanning-tree vlan** *vlan-id* global configuration command to enable STP on the desired VLAN.

Caution Switches that are not running spanning tree still forward BPDUs that they receive so that the other switches on the VLAN that have a running spanning-tree instance can break loops. Therefore, STP must be running on enough switches to break all the loops in the network; for example, at least one switch on each loop in the VLAN must be running STP. It is not absolutely necessary to run STP on all switches in the VLAN; however, if you are running STP only on a minimal set of switches, an incautious change to the network that introduces another loop into the VLAN can result in a broadcast storm.

Note If you have already used all available spanning-tree instances on your switch, adding another VLAN anywhere in the VTP domain creates a VLAN that is not running STP on that switch. If you have the default allowed list on the trunk ports of that switch, the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that will not be broken, particularly if there are several adjacent switches that have all run out of spanning-tree instances. You can prevent this possibility by setting allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances. Setting up allowed lists is not necessary in many cases and can make it more labor-intensive to add another VLAN to the network.

Spanning-tree commands determine the configuration of VLAN spanning-tree instances. You create a spanning-tree instance when you assign an interface to a VLAN. The spanning-tree instance is removed when the last interface is moved to another VLAN. You can configure switch and port parameters before a spanning-tree instance is created; these parameters are applied when the spanning-tree instance is created.

STP Overview

STP is a Layer 2 link management protocol that provides path redundancy while preventing undesirable loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. STP operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

When you create fault-tolerant internetworks, you must have a loop-free path between all nodes in a network. The spanning-tree algorithm calculates the best loop-free path throughout a switched Layer 2 network. Switches send and receive STP frames at regular intervals. The switches do not forward these frames, but use the frames to construct a loop-free path.

Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Switches might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an unstable network.

STP defines a tree with a root switch and a loop-free path from the root to all switches in the Layer 2 network. STP forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the spanning-tree algorithm recalculates the spanning-tree topology and activates the standby path.

When two interfaces on a switch are part of a loop, the STP port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. The STP port priority value represents the location of an interface in the network topology and how well it is located to pass traffic. The STP path cost value represents media speed.

Bridge ID, Switch Priority, and Extended System ID

The IEEE 802.1D standard requires that each switch has an unique bridge identifier (bridge ID), which determines the selection of the root switch. Because each VLAN is considered as a different *logical bridge* with PVST+, the same switch must have as many different bridge IDs as VLANs configured on it. Each VLAN on the switch has a unique 8-byte bridge ID; the two most-significant bytes are used for the switch priority, and the remaining six bytes are derived from the switch MAC address.

In Release 12.1(8)EA1 and later, Catalyst 3550 switches support the 802.1T spanning-tree extensions, and some of the bits previously used for the switch priority are now used as the VLAN identifier. The result is that fewer MAC addresses are reserved for the switch, and a larger range of VLAN IDs can be supported, all while maintaining the uniqueness of the bridge ID. As shown in Table 10-1, the two bytes previously used for the switch priority are reallocated into a 4-bit priority value and a 12-bit extended system ID value equal to the VLAN ID. In earlier releases, the switch priority is a 16-bit value.

Table 10-1 Switch Priority Value and Extended System ID

Switch Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	**Bit 15**	**Bit 14**	**Bit 13**	**Bit 12**	**Bit 11**	**Bit 10**	**Bit 9**	**Bit 8**	**Bit 7**	**Bit 6**	**Bit 5**	**Bit 4**	**Bit 3**	**Bit 2**	**Bit 1**
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

STP uses the extended system ID, the switch priority, and the allocated STP MAC address to make the bridge ID unique for each VLAN. With earlier releases, STP uses one MAC address per VLAN to make the bridge ID unique for each VLAN.

Support for the extended system ID affects how you manually configure the root switch, the secondary root switch, and the switch priority of a VLAN. For more information, see the "Configuring the Root Switch" section on page 10-22, "Configuring a Secondary Root Switch" section on page 10-24, and "Configuring the Switch Priority of a VLAN" section on page 10-28.

Election of the Root Switch

All switches in the Layer 2 network participating in STP gather information about other switches in the network through an exchange of data messages called Bridge Protocol Data Units (BPDUs). This exchange of messages results in these actions:

- The election of a unique root switch for each spanning-tree instance
- The election of a designated switch for every switched LAN segment

- The removal of loops in the switched network by blocking Layer 2 interfaces connected to redundant links

For each VLAN, the switch with the highest switch priority (the lowest numerical priority value) is elected as the root switch. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root switch. The switch priority value occupies the most significant bits of the bridge ID.

When you change the switch priority value, you change the probability that the switch will be elected as the root switch. Configuring a higher value increases the probability; a lower value decreases the probability.

The root switch is the logical center of the STP topology in a switched network. All paths that are not needed to reach the root switch from anywhere in the switched network are placed in STP blocking mode.

BPDUs contain information about the sending switch and its ports, including switch and MAC addresses, switch priority, port priority, and path cost. STP uses this information to elect the root switch and root port for the switched network and the root port and designated port for each switched segment.

Bridge Protocol Data Units

The stable, active STP topology of a switched network is determined by these elements:

- The unique bridge ID (switch priority and MAC address) associated with each VLAN on each switch
- The STP path cost to the root switch
- The port identifier (port priority and MAC address) associated with each Layer 2 interface

The BPDUs are sent in one direction from the root switch, and each switch sends configuration BPDUs to communicate and to compute the STP topology. Each configuration BPDU contains this information:

- The unique bridge ID of the switch that the sending switch identifies as the root switch
- The STP path cost to the root
- The bridge ID of the sending switch
- Message age
- The identifier of the sending interface
- Values for the hello, forward delay, and max-age protocol timers

When a switch sends a BPDU frame, all switches connected to the LAN on which the frame is sent receive the BPDU. When a switch receives a BPDU, it does not forward the frame but instead uses the information in the frame to calculate a BPDU, and, if the topology changes, starts a BPDU transmission.

A BPDU exchange results in these actions:

- One switch is elected as the root switch.
- The shortest distance to the root switch is calculated for each switch based on the path cost.
- A designated switch for each LAN segment is selected. The designated switch is the one closest to the root switch through which frames are forwarded to the root.
- A root port is selected. This port provides the best path from the switch to the root switch.
- Interfaces included in the spanning-tree instance are selected.
- All interfaces not included in the spanning tree are blocked.

STP Timers

Table 10-2 describes the STP timers that affect the entire spanning-tree performance.

Table 10-2 Spanning Tree Protocol Timers

Variable	Description
Hello timer	Determines how often the switch broadcasts hello messages to other switches.
Forward-delay timer	Determines how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Determines the amount of time the switch stores protocol information received on an interface.

Creating the STP Topology

In Figure 10-1, Switch A is elected as the root switch because the switch priority of all the switches is set to the default (32768) and Switch A has the lowest MAC address. However, because of traffic patterns, number of forwarding interfaces, or link types, Switch A might not be the ideal root switch. By increasing the priority (lowering the numerical value) of the ideal switch so that it becomes the root switch, you force an STP recalculation to form a new topology with the ideal switch as the root.

Figure 10-1 STP Topology

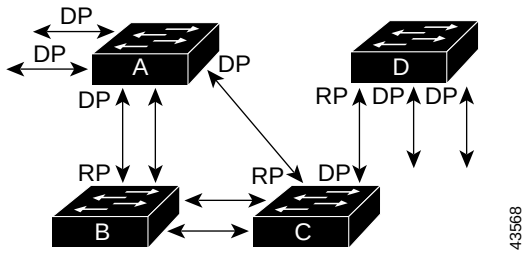

RP = Root Port
DP = Designated Port

When the spanning-tree topology is calculated based on default parameters, the path between source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to an interface that has a higher number than the current root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on Switch B is a Gigabit Ethernet link, and that another port on Switch B (a 10/100 link) is the root port. Network traffic might be more efficient over the Gigabit Ethernet link. By changing the STP port priority on the Gigabit Ethernet interface to a higher priority (lower numerical value) than the root port, the Gigabit Ethernet interface becomes the new root port.

STP Interface States

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When a Layer 2 interface transitions directly from nonparticipation in the spanning-tree topology to the forwarding state, it can create temporary data loops. Interfaces must wait for new topology information to propagate through the switched LAN before starting to forward frames. They must allow the frame lifetime to expire for forwarded frames that have used the old topology.

Each Layer 2 interface on a switch using STP exists in one of these states:

- Blocking—The Layer 2 interface does not participate in frame forwarding.
- Listening—the first transitional state after the blocking state when STP determines that the Layer 2 interface should participate in frame forwarding.
- Learning—The Layer 2 interface prepares to participate in frame forwarding.
- Forwarding—The Layer 2 interface forwards frames.
- Disabled—The Layer 2 interface is not participating in STP because of a shutdown port, no link on the port, or no spanning-tree instance running on the port.

A Layer 2 interface moves through these states:

- From initialization to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled

Figure 10-2 illustrates how a Layer 2 interface moves through the states.

Figure 10-2 Spanning Tree Layer 2 Interface States

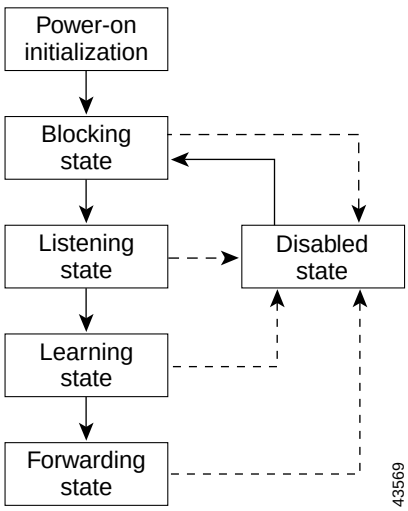

When you power up the switch, STP is enabled by default, and every interface in the switch, VLAN, or network goes through the blocking state and the transitory states of listening and learning. Spanning tree stabilizes each Layer 2 interface at the forwarding or blocking state.

When the spanning-tree algorithm places a Layer 2 interface in the forwarding state, this process occurs:

1. The Layer 2 interface is in the listening state while spanning tree waits for protocol information to transition the interface to the blocking state.
2. While spanning tree waits the forward-delay timer to expire, it moves the Layer 2 interface to the learning state and resets the forward-delay timer.
3. In the learning state, the Layer 2 interface continues to block frame forwarding as the switch learns end-station location information for the forwarding database.
4. When the forward-delay timer expires, spanning tree moves the Layer 2 interface to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 interface in the blocking state does not participate in frame forwarding. After initialization, a BPDU is sent to each Layer 2 interface in the switch. A switch initially functions as the root until it exchanges BPDUs with other switches. This exchange establishes which switch in the network is the root or root switch. If there is only one switch in the network, no exchange occurs, the forward-delay timer expires, and the interfaces move to the listening state. An interface always enters the blocking state after switch initialization.

A Layer 2 interface in the blocking state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Listening State

The listening state is the first state a Layer 2 interface enters after the blocking state. The Layer 2 interface enters this state when STP determines that the Layer 2 interface should participate in frame forwarding.

A Layer 2 interface in the listening state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Learning State

A Layer 2 interface in the learning state prepares to participate in frame forwarding. The Layer 2 interface enters the learning state from the listening state.

A Layer 2 interface in the learning state performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding

- Learns addresses
- Receives BPDUs

Forwarding State

A Layer 2 interface in the forwarding state forwards frames. The Layer 2 interface enters the forwarding state from the learning state.

A Layer 2 interface in the forwarding state performs as follows:

- Receives and forwards frames received on the port
- Forwards frames switched from another port
- Learns addresses
- Receives BPDUs

Disabled State

A Layer 2 interface in the disabled state does not participate in frame forwarding or STP. A Layer 2 interface in the disabled state is nonoperational.

A disabled Layer 2 interface performs as follows:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Does not receive BPDUs

STP Address Management

IEEE 802.1D specifies 17 multicast addresses, ranging from 0x00180C2000000 to 0x0180C2000010, to be used by different bridge protocols. These addresses are static addresses that cannot be removed.

Regardless of the STP state, the switch receives but does not forward packets destined for addresses between 0x0180c2000000 and 0x1080C200000F.

If STP is enabled, the switch CPU receives packets destined for 0x0180C2000000 and 0x0180C2000010. If STP is disabled, the switch forwards those packets as unknown multicast addresses.

STP and IEEE 802.1Q Trunks

The IEEE 802.1Q standard for VLAN trunks imposes some limitations on the spanning-tree strategy for a network. The standard requires only one spanning-tree instance for *all* VLANs allowed on the trunks. However, in a network of Cisco switches connected through 802.1Q trunks, the switches maintain one spanning-tree instance for *each* VLAN allowed on the trunks.

When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch uses per-VLAN spanning tree+ (PVST+) to provide STP interoperability. It combines the spanning-tree instance of the 802.1Q VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch.

However, all PVST+ information is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

PVST+ is automatically enabled on 802.1Q trunks, and no user configuration is required. The external spanning-tree behavior on access ports and Inter-Switch Link (ISL) trunks is not affected by PVST+.

For more information on 802.1Q trunks, see Chapter 9, "Creating and Maintaining VLANs."

VLAN-Bridge STP

Cisco VLAN-bridge STP is used with the fallback bridging feature (bridge groups), which forwards non-IP protocols such as DECnet or IPX between two or more VLAN bridge domains or routed ports. The VLAN-bridge STP allows the bridge groups to form a spanning tree on top of the individual VLAN spanning trees to prevent loops from forming if there are multiple connections among VLANs. It also prevents the individual spanning trees from the VLANs being bridged from collapsing into a single spanning tree.

To support VLAN-bridge STP, some of the spanning-tree timers are increased. For more information, see Chapter 26, "Configuring Fallback Bridging."

STP and Redundant Connectivity

You can create a redundant backbone with STP by connecting two switch interfaces to another device or to two different devices. STP automatically disables one interface but enables it if the other one fails, as shown in Figure 10-3. If one link is high-speed and the other is low-speed, the low-speed link is always disabled. If the speeds are the same, the port priority and port ID are added together, and STP disables the link with the lowest value.

Figure 10-3 STP and Redundant Connectivity

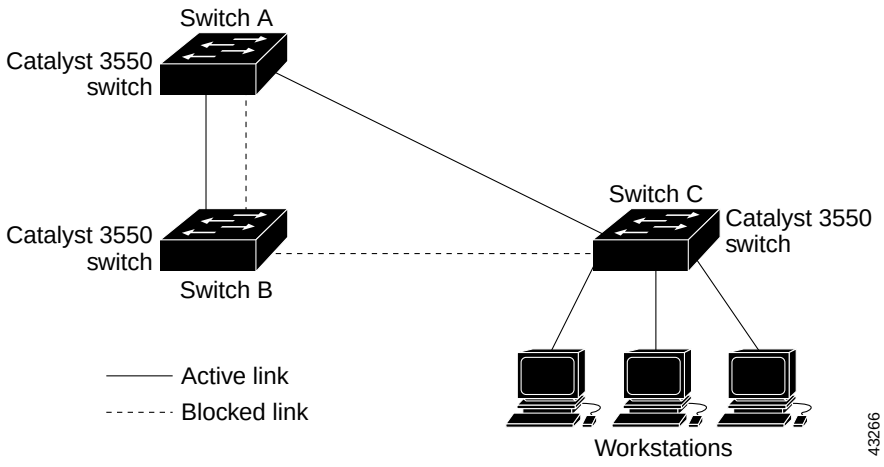

You can also create redundant links between switches by using EtherChannel groups. For more information, see the Chapter 21, "Configuring EtherChannel."

Accelerated Aging to Retain Connectivity

The default for aging dynamic addresses is 5 minutes, the default setting of the **mac-address-table aging-time** global configuration command. However, an STP reconfiguration can cause many station locations to change. Because these stations could be unreachable for 5 minutes or more during a reconfiguration, the address-aging time is accelerated so that station addresses can be dropped from the address table and then relearned. The accelerated aging is the same as the forward-delay parameter value (**spanning-tree vlan** *vlan-id* **forward-time** *seconds* global configuration command) when STP reconfigures.

Because each VLAN is a separate spanning-tree instance, the switch accelerates aging on a per-VLAN basis. An STP reconfiguration on one VLAN can cause the dynamic addresses learned on that VLAN to be subject to accelerated aging. Dynamic addresses on other VLANs can be unaffected and remain subject to the aging interval entered for the switch.

Understanding Advanced STP Features

This section describes how advanced STP features work. It includes this information:

- Understanding Port Fast, page 10-10
- Understanding BPDU Guard, page 10-11
- Understanding UplinkFast, page 10-12
- Understanding Cross-Stack UplinkFast, page 10-13
- Understanding BackboneFast, page 10-18
- Understanding Root Guard, page 10-20
- Understanding EtherChannel Guard, page 10-20

For configuration information, see the "Configuring Advanced STP Features" section on page 10-32.

Understanding Port Fast

Port Fast immediately brings an interface configured as a Layer 2 access port to the forwarding state from a blocking state, bypassing the listening and learning states. You can use Port Fast on Layer 2 access ports connected to a single workstation or server, as shown in Figure 10-4, to allow those devices to immediately connect to the network, rather than waiting for STP to converge.

If the interface receives a BPDU, which should not happen if the interface is connected to a single workstation or server, STP puts the port in the blocking state. An interface with Port Fast enabled goes through the normal cycle of STP status changes when the switch is restarted. For more information, see the "Configuring Port Fast" section on page 10-32.

Note Because the purpose of Port Fast is to minimize the time access ports must wait for STP to converge, it is effective only when used on access ports. If you enable Port Fast on a port connecting to another switch, you risk creating a spanning-tree loop.

Figure 10-4 Port Fast-Enabled Ports

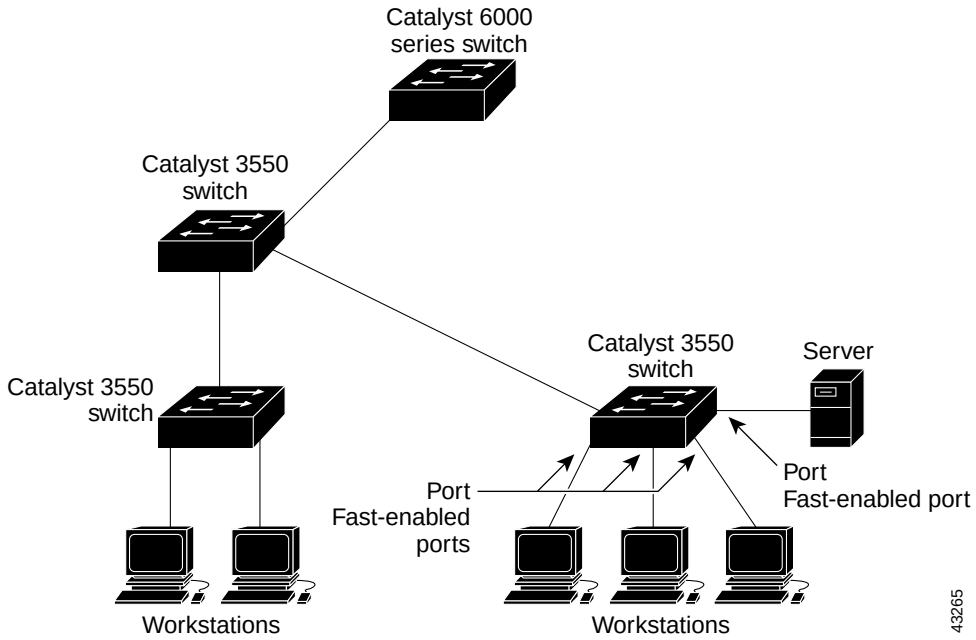

Understanding BPDU Guard

When the BPDU guard feature is enabled on the switch, STP shuts down Port Fast-enabled interfaces that receive BPDUs rather than putting them into the blocking state. In a valid configuration, Port Fast-enabled interfaces do not receive BPDUs. Receipt of a BPDU by a Port Fast-enabled interface means an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature places the interface into the ErrDisable state. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface back in service. For more information, see the "Configuring BPDU Guard" section on page 10-33.

Note When enabled on the switch, STP applies the BPDU guard feature to all Port Fast-enabled interfaces.

Understanding UplinkFast

Switches in hierarchical networks can be grouped into backbone switches, distribution switches, and access switches. Figure 10-5 shows a complex network where distribution switches and access switches each have at least one redundant link that STP blocks to prevent loops.

Figure 10-5 Switches in a Hierarchical Network

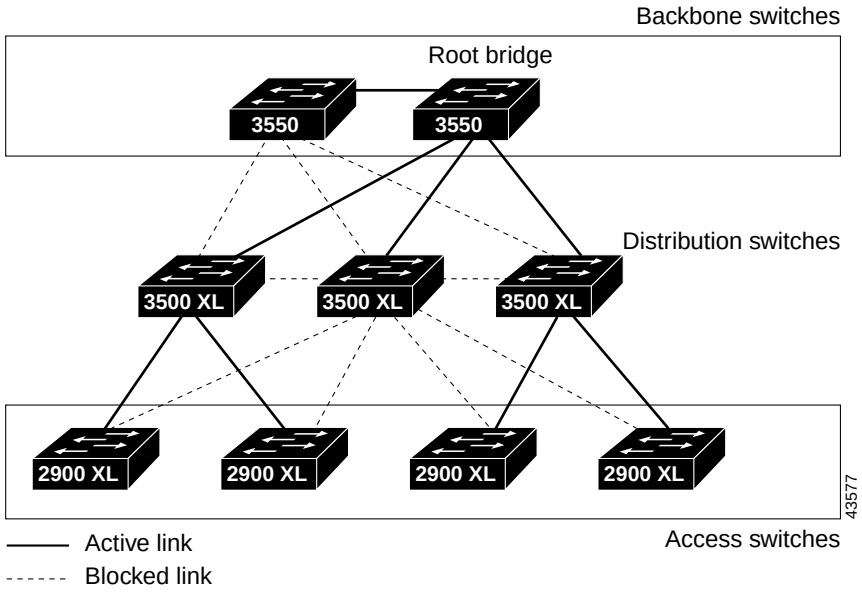

If a switch looses connectivity, it begins using the alternate paths as soon as STP selects a new root port. By using STP UplinkFast, you can accelerate the choice of a new root port when a link or switch fails or when STP reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with normal STP procedures.

When STP reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface. You can limit these bursts of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the STP topology converges more slowly after a loss of connectivity.

Note UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

Figure 10-6 shows an example topology with no link failures. Switch A, the root switch, is connected directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that is connected directly to Switch B is in a blocking state.

Figure 10-6 UplinkFast Example Before Direct Link Failure

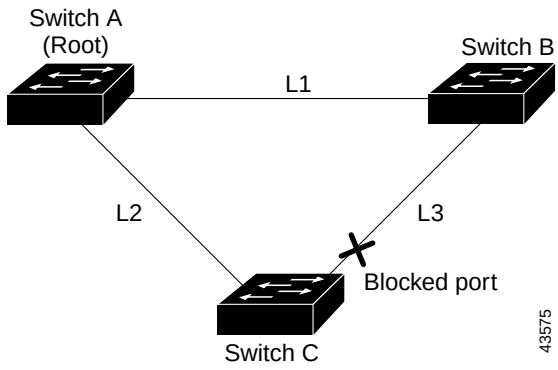

If Switch C detects a link failure on the currently active link L2 on the root port (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 10-7. This change takes approximately 1 to 5 seconds. For more information, see the "Configuring UplinkFast for Use with Redundant Links" section on page 10-34.

Figure 10-7 UplinkFast Example After Direct Link Failure

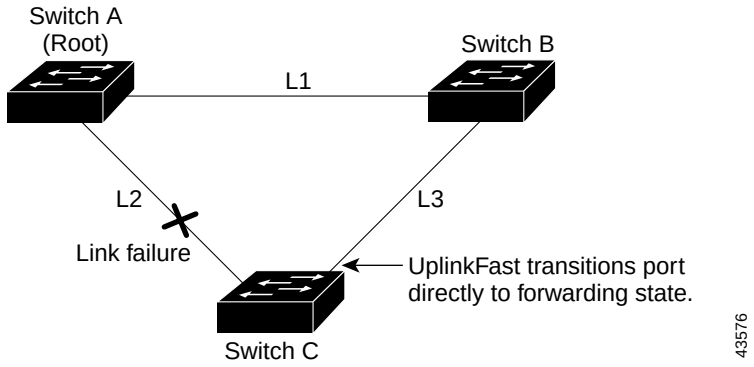

Understanding Cross-Stack UplinkFast

Cross-stack UplinkFast (CSUF) provides a fast spanning-tree transition (fast convergence in less than 1 second under normal network conditions) across a stack of switches that use the GigaStack GBICs connected in a shared cascaded configuration (multidrop backbone). During the fast transition, an alternate redundant link on the stack of switches is placed in the forwarding state without causing temporary spanning-tree loops or loss of connectivity to the backbone. With this feature, you can have a redundant and resilient network in some configurations.

CSUF might not provide a fast transition all the time; in these cases, the normal STP transition occurs, completing in 30 to 40 seconds. For more information, see the "Events that Cause Fast Convergence" section on page 10-15. For configuration information, see the "Configuring Cross-Stack UplinkFast" section on page 10-35.

How CSUF Works

CSUF ensures that one link in the stack is elected as the path to the root. As shown in Figure 10-8, Switches A, B, and C are cascaded through the GigaStack GBIC to form a multidrop backbone, which communicates control and data traffic across the switches at the access layer. The switches in the stack use their stack ports to communicate with each other and to connect to the stack backbone; stack ports are always in the STP forwarding state. The stack-root port on Switch A provides the path to the root of the spanning tree; the alternate stack-root ports on Switches B and C can provide an alternate path to the spanning-tree root if the current stack-root switch fails or if its link to the spanning-tree root fails.

Link A, the root link, is in the STP forwarding state; Links B and C are alternate redundant links that are in the STP blocking state. If Switch A fails, if its stack-root port fails, or if Link A fails, CSUF selects either the Switch B or Switch C alternate stack-root port and puts it into the forwarding state in less than 1 second.

Figure 10-8 Cross-Stack UplinkFast Topology

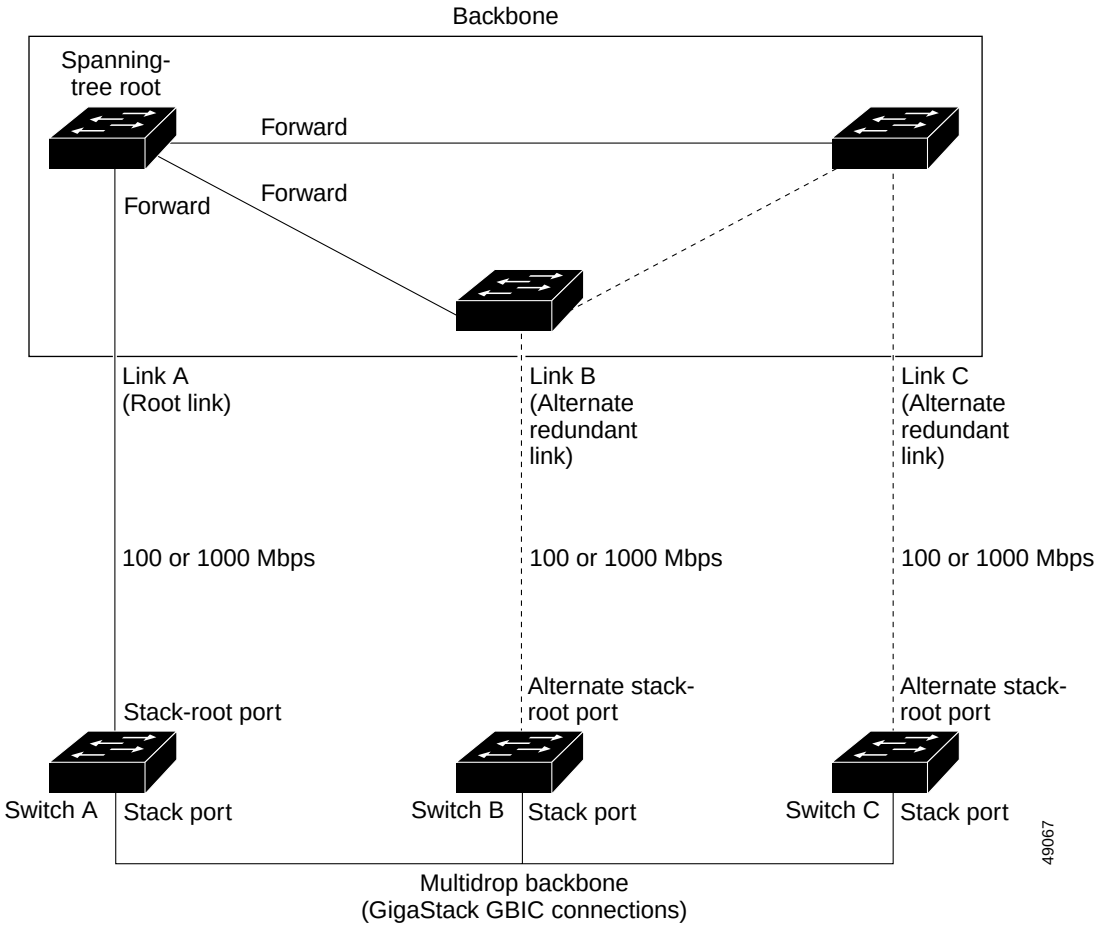

CSUF implements the Stack Membership Discovery Protocol and the Fast Uplink Transition Protocol. Using the Stack Membership Discovery Protocol, all stack switches build a neighbor list of stack members through the receipt of discovery hello packets. When certain link loss or STP events occur (described in "Events that Cause Fast Convergence" section on page 10-15), the Fast Uplink Transition Protocol uses the neighbor list to send fast-transition requests on the stack port to stack members.

The switch sending the fast-transition request needs to do a fast transition to the forwarding state of a port that it has chosen as the root port, and it must obtain an acknowledgement from each stack switch before performing the fast transition.

Each switch in the stack determines if the sending switch is a better choice than itself to be the stack root of this spanning-tree instance by comparing STP root, cost, and bridge ID. If the sending switch is the best choice as the stack root, each switch in the stack returns an acknowledgement; otherwise, it does not respond to the sending switch (drops the packet). The sending switch then has not received acknowledgements from all stack switches.

When acknowledgements are received from all stack switches, the Fast Uplink Transition Protocol on the sending switch immediately transitions its alternate stack-root port to the forwarding state. If acknowledgements from all stack switches are not obtained by the sending switch, the normal STP transitions (blocking, listening, learning, and forwarding) take place, and the spanning-tree topology converges at its normal rate (2 * forward-delay time + max-age time).

The Fast Uplink Transition Protocol is implemented on a per-VLAN basis and affects only one spanning-tree instance at a time.

Events that Cause Fast Convergence

Depending on the network event or failure, the CSUF fast convergence might or might not occur.

Fast convergence (less than 1 second under normal network conditions) occurs under these circumstances:

- The stack-root port link fails.

  If two switches in the stack have alternate paths to the root, only one of the switches performs the fast transition.
- The failed link, which connects the stack root to the STP root, recovers.
- A network reconfiguration causes a new stack-root switch to be selected.
- A network reconfiguration causes a new port on the current stack-root switch to be chosen as the stack-root port.

Note The fast transition might not occur if multiple events occur simultaneously. For example, if a stack member switch is powered off, and at the same time, the link connecting the stack root to the STP root comes back up, the normal STP convergence occurs.

Normal STP convergence (30 to 40 seconds) occurs under these conditions:

- The stack-root switch is powered off, or the software failed.
- The stack-root switch, which was powered off or failed, is powered on.
- A new switch, which might become the stack root, is added to the stack.
- A switch other than the stack root is powered off or failed.
- A link fails between stack ports on the multidrop backbone.

Limitations

These limitations apply to CSUF:

- CSUF uses the GigaStack GBIC and runs on all Catalyst 3550 switches, all Catalyst 3500 XL switches, but only on modular Catalyst 2900 XL switches that have the 1000BASE-X module installed.
- Up to nine stack switches can be connected through their stack ports to the multidrop backbone. Only one stack port per switch is supported.
- Each stack switch can be connected to the STP backbone through one uplink.
- If the stack consists of a mixture of Catalyst 3550, Catalyst 3500 XL, and Catalyst 2900 XL switches, up to 64 VLANs with STP enabled are supported. If the stack consists of Catalyst 3550 switches, up to 128 VLANs with STP enabled are supported.

Connecting the Stack Ports

A fast transition occurs across the stack of switches if the multidrop backbone connections are a continuous link from one GigaStack GBIC to another as shown in the top half of Figure 10-9. The bottom half of Figure 10-9 shows how to connect the GigaStack GBIC to achieve a normal convergence time.

You should follow these guidelines:

- A switch supports only one stack port.
- Do not connect alternate stack-root ports to stack ports.
- Connect all stack ports on the switch stack to the multidrop backbone.
- You can connect the open ports on the top and bottom GigaStack GBICs within the same stack to form a redundant link.

Figure 10-9 GigaStack GBIC Connections and STP Convergence

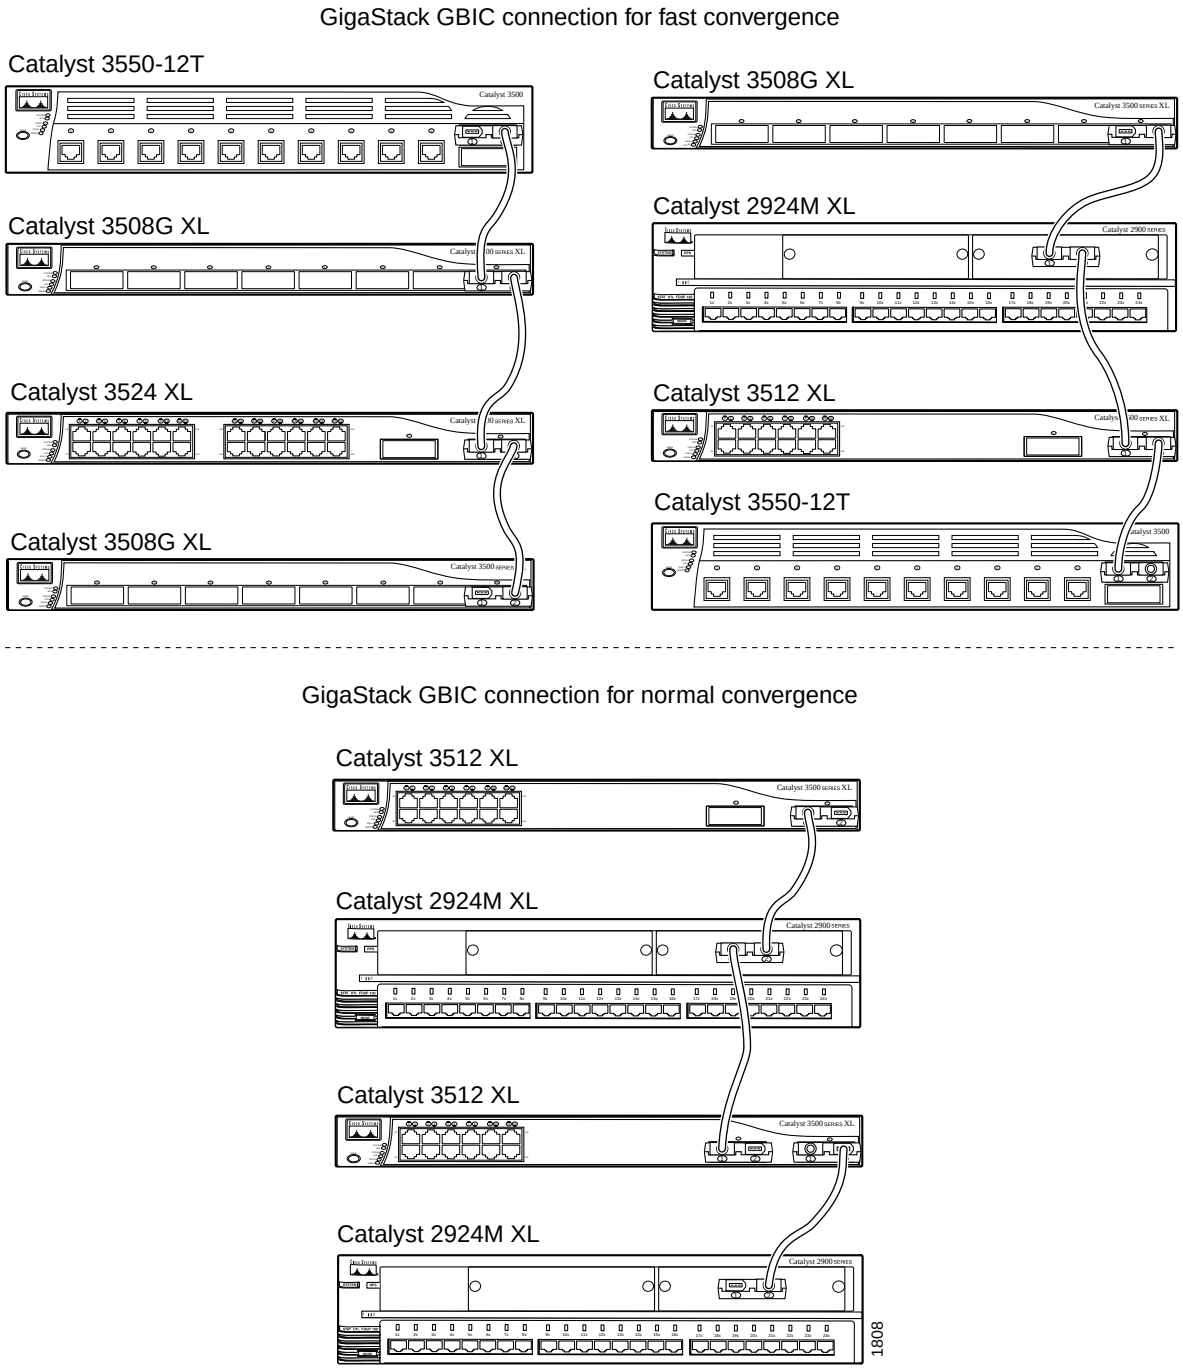

Understanding BackboneFast

BackboneFast is started when a root port or blocked port on a switch receives inferior BPDUs from its designated bridge. An inferior BPDU identifies one switch as both the root bridge and the designated bridge. When a switch receives an inferior BPDU, it means that a link to which the switch is not directly connected (an *indirect* link) has failed (that is, the designated bridge has lost its connection to the root switch). Under STP rules, the switch ignores inferior BPDUs for the configured maximum aging time specified by the **spanning-tree max-age** global configuration command.

The switch tries to determine if it has an alternate path to the root switch. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root switch. (Self-looped ports are not considered alternate paths to the root switch.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root switch, causes the maximum aging time on the root to expire, and becomes the root switch according to normal STP rules.

If the switch has alternate paths to the root switch, it uses these alternate paths to send a new kind of Protocol Data Unit (PDU) called the Root Link Query PDU. The switch sends the Root Link Query PDU on all alternate paths to the root switch. If the switch determines that it still has an alternate path to the root, it causes the maximum aging time on the ports on which it received the inferior BPDU to expire. If all the alternate paths to the root switch indicate that the switch has lost connectivity to the root switch, the switch causes the maximum aging times on the ports on which it received an inferior BPDU to expire. If one or more alternate paths can still connect to the root switch, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them out of the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.

Figure 10-10 shows an example topology with no link failures. Switch A, the root switch, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that connects directly to Switch B is in the blocking state.

Figure 10-10 BackboneFast Example Before Indirect Link Failure

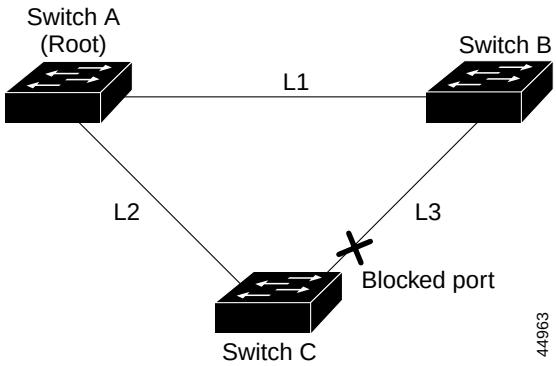

If link L1 fails, Switch C cannot detect this failure because it is not connected directly to link L1. However, because Switch B is directly connected to the root switch over L1, it detects the failure, elects itself the root, and begins sending BPDUs to Switch C, identifying itself as the root. When Switch C receives the inferior BPDUs from Switch B, Switch C assumes that an indirect failure has occurred. At that point, BackboneFast allows the blocked port on Switch C to move immediately to the listening state without waiting for the maximum aging time for the port to expire. BackboneFast then transitions the Layer 2 interface on Switch C to the forwarding state, providing a path from Switch B to Switch A. This

switchover takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. Figure 10-11 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 10-11 BackboneFast Example After Indirect Link Failure

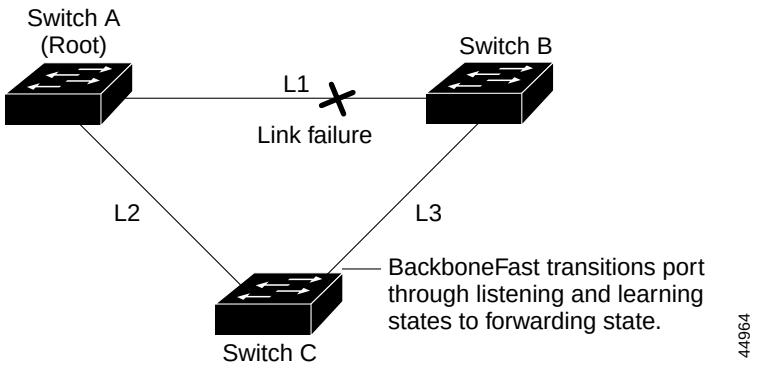

If a new switch is introduced into a shared-medium topology as shown in Figure 10-12, BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated bridge (Switch B). The new switch begins sending inferior BPDUs that say it is the root switch. However, the other switches ignore these inferior BPDUs, and the new switch learns that Switch B is the designated bridge to Switch A, the root switch. For more information, see the "Configuring BackboneFast" section on page 10-36.

Figure 10-12 Adding a Switch in a Shared-Medium Topology

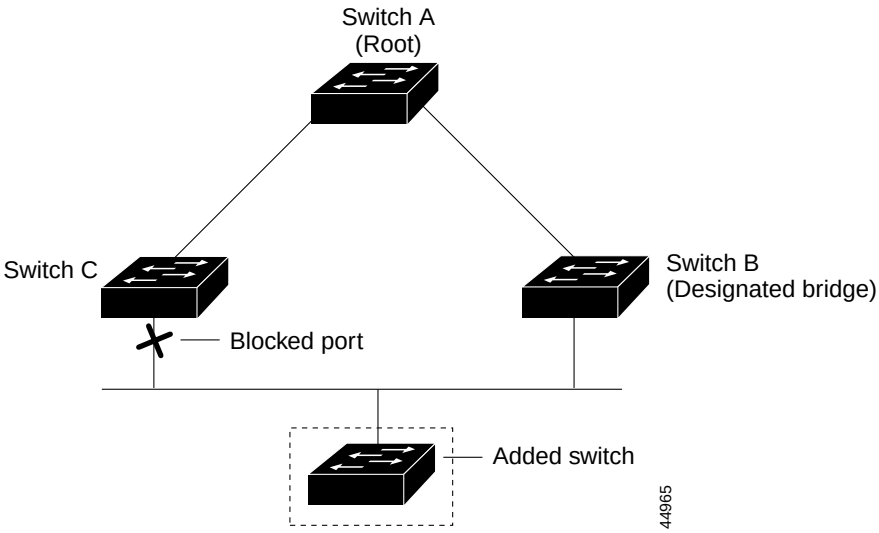

Understanding Root Guard

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, STP can reconfigure itself and select a *customer switch* as the STP root switch, as shown in Figure 10-13. You can avoid this situation by configuring the root-guard feature on interfaces that connect to switches outside of your customer's network. If STP calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the customer's switch from becoming the root switch or being in the path to the root.

If a switch outside the network becomes the root switch, the interface is blocked (root-inconsistent state), and STP selects a new root switch. The customer's switch does not become the root switch and is not in the path to the root. For more information, see the "Configuring Root Guard" section on page 10-36.

Caution Misuse of the root-guard feature can cause a loss of connectivity.

Figure 10-13 STP in a Service-Provider Network

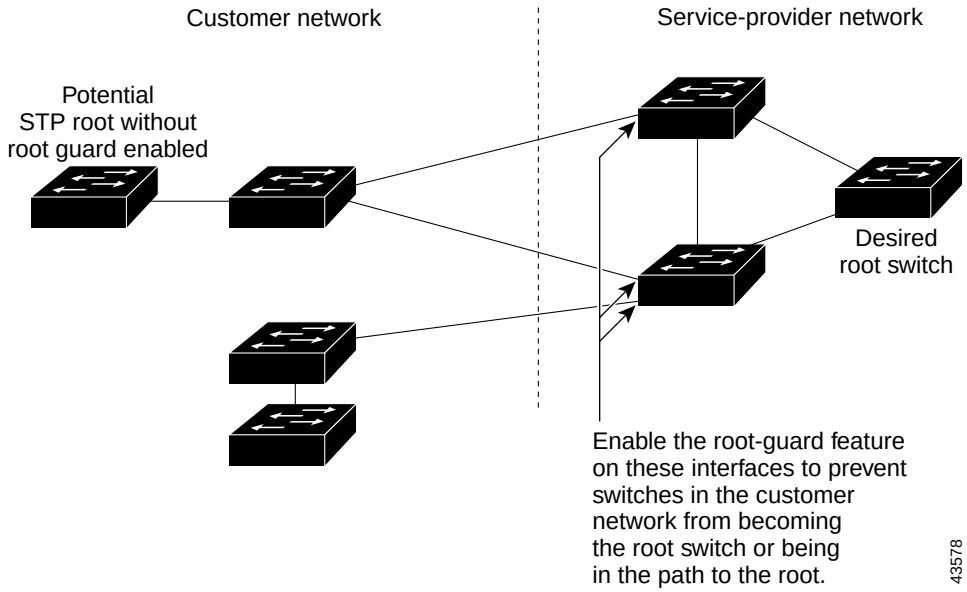

Understanding EtherChannel Guard

EtherChannel guard detects a misconfigured EtherChannel when Catalyst 3550 switch interfaces are configured as an EtherChannel while interfaces on the other device are not or not all the interfaces on the other device are in the same EtherChannel. This feature is enabled by default.

In response to misconfiguration detected on the other device, EtherChannel guard puts Catalyst 3550 interfaces into the error-disabled (err-disabled) state to prevent a spanning-tree loop. For more information, see the "Enabling EtherChannel Guard" section on page 10-37.

Configuring Basic STP Features

These sections include basic STP configuration information:

- Default STP Configuration, page 10-21
- Disabling STP, page 10-22
- Configuring the Root Switch, page 10-22
- Configuring a Secondary Root Switch, page 10-24
- Configuring STP Port Priority, page 10-26
- Configuring STP Path Cost, page 10-27
- Configuring the Switch Priority of a VLAN, page 10-28
- Configuring the Hello Time, page 10-29
- Configuring the Forwarding-Delay Time for a VLAN, page 10-29
- Configuring the Maximum-Aging Time for a VLAN, page 10-30
- Configuring STP for Use in a Cascaded Stack, page 10-30
- Displaying STP Status, page 10-31

For advanced configuration information, see the "Configuring Advanced STP Features" section on page 10-32.

Default STP Configuration

Table 10-3 shows the default STP configuration.

Table 10-3 Default STP Configuration

Feature	Default Setting
Enable state	Enabled on VLAN 1. Up to 128 spanning-tree instances can be enabled.
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis—used on interfaces configured as Layer 2 access ports)	128.
Spanning-tree port cost (configurable on a per-interface basis—used on interfaces configured as Layer 2 access ports)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Spanning-tree VLAN port priority (configurable on a per-VLAN basis—used on interfaces configured as Layer 2 trunk ports)	128.
Spanning-tree VLAN port cost (configurable on a per-VLAN basis—used on interfaces configured as Layer 2 trunk ports)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Hello time	2 seconds.

Table 10-3 Default STP Configuration (continued)

Feature	Default Setting
Forward-delay time	15 seconds.
Maximum-aging time	20 seconds.
Port Fast	Disabled on all interfaces.
BPDU guard	Disabled on the switch.
UplinkFast	Disabled on the switch.
BackboneFast	Disabled on the switch.
Root guard	Disabled on all interfaces.
EtherChannel guard	Enabled on the switch.

Disabling STP

STP is enabled by default on VLAN 1 and on all newly created VLANs up to the spanning-tree limit specified in Table 10-3. Disable STP only if you are sure there are no loops in the network topology.

Caution When STP is disabled and loops are present in the topology, excessive traffic and indefinite packet duplication can drastically reduce network performance.

Beginning in privileged EXEC mode, follow these steps to disable STP on a per-VLAN basis:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no spanning-tree vlan** *vlan-id*	Disable STP on a per-VLAN basis. For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id*	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To re-enable STP, use the **spanning-tree vlan** *vlan-id* global configuration command.

Configuring the Root Switch

The switch maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.

To configure a switch to become the root, use the **spanning-tree vlan** *vlan-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value so that the switch becomes the root switch for the specified VLAN. When you enter this command, the switch checks the switch priority of the current root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.

If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in Table 10-1 on page 10-3.)

Note The **spanning-tree vlan** *vlan-id* **root** global configuration command fails if the value necessary to be the root switch is less than 1.

Before Release 12.1(8)EA1, entering the **spanning-tree vlan** *vlan-id* **root** global configuration command on a a Catalyst 3550 switch (no extended system ID) causes it to set its own switch priority for the specified VLAN to 8192 if this value causes this switch to become the root for the specified VLAN. If any root switch for the specified VLAN has a switch priority lower than 8192, the switch sets its own priority for the specified VLAN to 1 less than the lowest switch priority.

These examples show the effect of the **spanning-tree vlan** *vlan-id* **root** command with and without the extended system ID support:

- For Catalyst 3550 switch with the extended system ID (Release 12.1(8)EA1 and later), if all network devices in VLAN 20 have the default priority of 32768, entering the **spanning-tree vlan 20 root primary** command on the switch sets the switch priority to 24576, which causes this switch to become the switch bridge for VLAN 20.
- For Catalyst 3550 switches without the extended system ID (software before Release 12.1(8)EA1), if all network devices in VLAN 100 have the default priority of 32768, entering the **spanning-tree vlan 100 root primary** command on the switch sets the switch priority for VLAN 100 to 8192, which causes this switch to become the root switch for VLAN 100.

Note If your network consists of Catalyst 3550 switches that do not support the extended system ID and Catalyst 3550 switches that do support it, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

Note The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

Note We recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time after configuring the switch as the root switch.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the root switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **root primary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configure a switch as the root switch. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • (Optional) For **diameter** *net-diameter*, specify the maximum number of switches between any two end stations. The range is 2 to 7. • (Optional) For **hello-time** *seconds*, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **root** global configuration command.

Configuring a Secondary Root Switch

When you configure a Catalyst 3550 switch that supports the extended system ID as the secondary root, the STP switch priority is modified from the default value (32768) to 28672 so that the switch is likely to become the root switch for the specified VLAN if the primary root switch fails (if the other switches in the network use the default switch priority of 32768, and therefore, are unlikely to become the root switch). For Catalyst 3550 switches without the extended system ID support (software before Release 12.1(8)EA1), the switch priority is changed to 16384.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values as you used when configuring the primary root switch.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the secondary root switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **root secondary** [**diameter** *net-diameter* [**hello-time** *seconds*]]	Configure a switch as the secondary root switch. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • (Optional) For **diameter** *net-diameter*, specify the maximum number of switches between any two end stations. The range is 2 to 7. • (Optional) For **hello-time** *seconds*, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. Use the same network diameter and hello-time values that you used when configuring the primary root switch.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **root** global configuration command.

Configuring STP Port Priority

If a loop occurs, STP uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, STP puts the interface with the lowest interface number in the forwarding state and blocks other interfaces. The priority range is 0 to 255; the default is 128.

Cisco IOS uses the port priority value when the interface is configured as an access port and uses VLAN port priority values when the interface is configured as a trunk port.

Beginning in privileged EXEC mode, follow these steps to configure the STP port priority of an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
Step 3	**spanning-tree port-priority** *priority*	Configure the port priority for an interface that is an access port. For *priority*, the range is 0 to 255; the default is 128. The lower the number, the higher the priority.
Step 4	**spanning-tree vlan** *vlan-id* **port-priority** *priority*	Configure the VLAN port priority for an interface that is a trunk port. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *priority*, the range is 0 to 255; the default is 128. The lower the number, the higher the priority.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show spanning-tree interface** *interface-id* or **show spanning-tree vlan** *vlan-id*	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only if the port is in a link-up operative state. Otherwise, you can use the **show running-config interface** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree vlan** *vlan-id* **port-priority** interface configuration command. For information on how to configure load sharing on trunk ports by using STP port priorities, see the "Load Sharing Using STP" section on page 9-29.

Configuring STP Path Cost

The STP path cost default value is derived from the media speed of an interface. If a loop occurs, STP uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, STP puts the interface with the lowest interface number in the forwarding state and blocks other interfaces.

STP uses the cost value when the interface is configured as an access port and uses VLAN port cost values when the interface is configured as a trunk port.

Beginning in privileged EXEC mode, follow these steps to configure the STP cost of an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
Step 3	**spanning-tree cost** *cost*	Configure the cost for an interface that is an access port. If a loop occurs, STP uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. For *cost*, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 4	**spanning-tree vlan** *vlan-id* **cost** *cost*	Configure the VLAN cost for an interface that is a trunk port. If a loop occurs, STP uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *cost*, the range is 1 to 65535; the default value is derived from the media speed of the interface.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show spanning-tree interface** *interface-id* or **show spanning-tree vlan** *vlan-id*	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only for ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree cost** or the **no spanning-tree vlan** *vlan-id* **cost** interface configuration command. For information on how to configure load sharing on trunk ports using STP path costs, see the "Load Sharing Using STP" section on page 9-29.

Configuring the Switch Priority of a VLAN

You can configure the switch priority and make it more likely that the switch will be chosen as the root switch.

Note Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan** *vlan-id* **root primary** and the **spanning-tree vlan** *vlan-id* **root secondary** global configuration commands to modify the switch priority.

Beginning in privileged EXEC mode, follow these steps to configure the STP switch priority of a VLAN:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **priority** *priority*	Configure the switch priority of a VLAN. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *priority*, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Valid priority values are 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id* **bridge** [**brief**]	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **priority** global configuration command.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root switch by changing the STP hello time.

Note Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan** *vlan-id* **root primary** and the **spanning-tree vlan** *vlan-id* **root secondary** global configuration commands to modify the hello time.

Beginning in privileged EXEC mode, follow these steps to configure the STP hello time of a VLAN:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **hello-time** *seconds*	Configure the hello time of a VLAN. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *seconds*, the range is 1 to 10 seconds; the default is 2 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id* **bridge** [**brief**]	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **hello-time** global configuration command.

Configuring the Forwarding-Delay Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the STP forwarding-delay time for a VLAN:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **forward-time** *seconds*	Configure the forward time of a VLAN. The forward delay is the number of seconds a port waits before changing from its STP learning and listening states to the forwarding state. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *seconds*, the range is 4 to 30 seconds; the default is 15 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id* **bridge** [**brief**]	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **forward-time** global configuration command.

Configuring the Maximum-Aging Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the STP maximum-aging time for a VLAN:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree vlan** *vlan-id* **max-age** *seconds*	Configure the maximum-aging time of a VLAN. The maximum-aging time is the number of seconds a switch waits without receiving STP configuration messages before attempting a reconfiguration. • For *vlan-id*, the range is 1 to 1005. Do not enter leading zeros. • For *seconds*, the range is 6 to 40 seconds; the default is 20 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id* **bridge** [**brief**]	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree vlan** *vlan-id* **max-age** global configuration command.

Configuring STP for Use in a Cascaded Stack

STP uses default values that can be reduced when configuring your switch in cascaded configurations. If an STP root switch is part of a cluster that is one switch from a cascaded stack, you can customize STP to reconverge more quickly after a switch failure. Figure 10-14 shows switches in three cascaded stacks that use the GigaStack GBIC. Table 10-4 shows the default STP settings and those that are acceptable for these configurations.

Table 10-4 Default and Acceptable STP Parameter Settings (in seconds)

STP Parameter	STP Default	Acceptable for Option 1	Acceptable for Option 2	Acceptable for Option 3
Hello Time	2	1	1	1
Max Age	20	6	10	6
Forwarding Delay	15	4	7	4

Figure 10-14 Gigabit Ethernet Stack

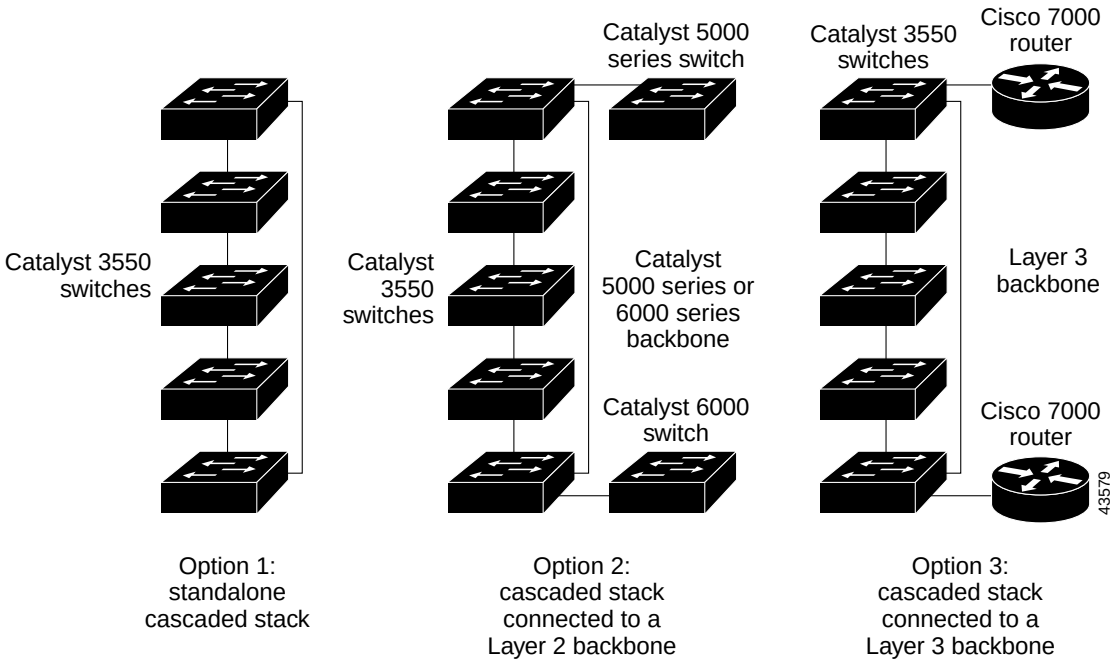
Catalyst 5000 series switch
Catalyst 3550 switches
Cisco 7000 router
Catalyst 3550 switches
Catalyst 5000 series or 6000 series backbone
Layer 3 backbone
Catalyst 6000 switch
Cisco 7000 router
43579
Option 1: standalone cascaded stack
Option 2: cascaded stack connected to a Layer 2 backbone
Option 3: cascaded stack connected to a Layer 3 backbone

Displaying STP Status

To display the current STP status, use one or more of the privileged EXEC commands in Table 10-5:

Table 10-5 Commands for Displaying STP Status

Command	Purpose
show spanning-tree active	Displays STP information on active interfaces only.
show spanning-tree brief	Displays a summary of interface information.
show spanning-tree interface *interface-id*	Displays information for the specified interface.
show spanning-tree summary [**totals**]	Displays a summary of port states or displays the total lines of the STP state section.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Configuring Advanced STP Features

These sections include advanced STP configuration information:

- Configuring Port Fast, page 10-32
- Configuring BPDU Guard, page 10-33
- Configuring UplinkFast for Use with Redundant Links, page 10-34
- Configuring Cross-Stack UplinkFast, page 10-35
- Configuring BackboneFast, page 10-36
- Configuring Root Guard, page 10-36
- Enabling EtherChannel Guard, page 10-37

Configuring Port Fast

A port with the Port Fast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.

Caution Use Port Fast *only* when connecting a single end station to a Layer 2 access port. Enabling this feature on an interface connected to a switch or hub could prevent STP from detecting and disabling loops in your network, which could cause broadcast storms and address-learning problems.

Beginning in privileged EXEC mode, follow these steps to enable Port Fast on a Layer 2 access port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
Step 3	**spanning-tree portfast**	Enable Port Fast on a Layer 2 access port connected to a single workstation or server. By default, Port Fast is disabled on all interfaces.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running interface** *interface-id*	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the Port Fast feature, use the **no spanning-tree portfast** interface configuration command.

Configuring BPDU Guard

When the BPDU guard feature is enabled on the switch, STP shuts down Port Fast-enabled interfaces that receive BPDUs rather than putting them into the blocking state.

Caution The BPDU guard feature works on Port Fast-enable interfaces. Configure Port Fast only on interfaces that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

Beginning in privileged EXEC mode, follow these steps to enable the BPDU guard feature on the switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree portfast bpduguard**	Enable BPDU guard on the switch. By default, BPDU guard is disabled on the switch.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree summary total**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

In a valid configuration, Port Fast-enabled interfaces do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled interface means an invalid configuration, such as the connection of an unauthorized device. If a BPDU is received on Port Fast-enabled interface, the BPDU guard feature places the interface into the ErrDisable state. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface back in service.

To disable BPDU guard, use the **no spanning-tree portfast bpduguard** global configuration command.

Configuring UplinkFast for Use with Redundant Links

UplinkFast increases the switch priority to 49152 and adds 3000 to the STP path cost only if the port used the default path cost before you enabled UplinkFast, making it unlikely that the switch will become the root switch. The **max-update-rate** represents the number of multicast packets sent per second (the default is 150 packets per second). UplinkFast cannot be enabled on VLANs that have been configured for switch priority. To enable UplinkFast on a VLAN with switch priority configured, first restore the switch priority on the VLAN to the default value by using a **no spanning-tree vlan** *vlan-id* **priority** global configuration command.

Note When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

Beginning in privileged EXEC mode, follow these steps to enable UplinkFast:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree uplinkfast** [**max-update-rate** *pkts-per-second*]	Enable UplinkFast on the switch. For *pkts-per-second*, the range is 0 to 65535 packets per second; the default is 150. If you set the rate to 0, station-learning frames are not generated, and the STP topology converges more slowly after a loss of connectivity.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

When UplinkFast is enabled, the switch priority of all VLANs is set to 49152, and the path cost of all interfaces and VLAN trunks is increased by 3000 if you did not modify the path cost from its default setting. This change reduces the chance that the switch will become the root port. When UplinkFast is disabled, the switch priorities of all VLANs and path costs of all interfaces are set to default values if you did not modify them from their defaults.

To return the update packet rate to the default setting, use the **no spanning-tree uplinkfast max-update-rate** global configuration command. To disable UplinkFast, use the **no spanning-tree uplinkfast** command.

Configuring Cross-Stack UplinkFast

Before enabling CSUF, make sure your stack switches are properly connected. For more information, see the "Connecting the Stack Ports" section on page 10-16.

Beginning in privileged EXEC mode, follow these steps to enable CSUF:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree uplinkfast** [**max-update-rate** *pkts-per-second*]	Enable UplinkFast on the switch. (Optional) For **max-update-rate** *pkts-per-second*, specify the number of packets per second at which update packets are sent. The range is 0 to 65535; the default is 150 packets per second.
Step 1	**interface** *interface-id*	Enter interface configuration mode, and specify the GBIC interface on which to enable CSUF.
Step 2	**spanning-tree stack-port**	Enable CSUF on only one stack-port GBIC interface. The stack port connects to the GigaStack GBIC multidrop backbone. If you try to enable CSUF on a Fast Ethernet or a Gigabit-capable Ethernet port, you receive an error message. If CSUF is already enabled on an interface and you try to enable it on another interface, you receive an error message. You must disable CSUF on the first interface before enabling it on a new interface. Use this command only on access switches.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable CSUF on an interface, use the **no spanning-tree stack-port** interface configuration command. To disable UplinkFast on the switch and all its VLANs, use the **no spanning-tree uplinkfast** global configuration command.

Configuring BackboneFast

You can enable BackboneFast to detect indirect link failures and to start the spanning-tree reconfiguration sooner.

Note If you use BackboneFast, you must enable it on all switches in the network. BackboneFast is not supported on Token Ring VLANs. This feature is supported for use with third-party switches.

Beginning in privileged EXEC mode, follow these steps to enable BackboneFast:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree backbonefast**	Enable BackboneFast on the switch.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree vlan** *vlan-id*	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the BackboneFast feature, use the **no spanning-tree backbonefast** global configuration command.

Configuring Root Guard

Root guard enabled on an interface applies to all the VLANs to which the interface belongs. Each VLAN has its own spanning-tree instance.

Do not enable the root guard on interfaces to be used by the UplinkFast feature. With UplinkFast, the backup interfaces (in the blocked state) replace the root port in the case of a failure. However, if root guard is also enabled, all the backup interfaces used by the UplinkFast feature are placed in the root-inconsistent state (blocked) and are prevented from reaching the forwarding state.

Beginning in privileged EXEC mode, follow these steps to enable root guard on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*).
Step 3	**spanning-tree guard root**	Enable root guard on the interface. By default, root guard is disabled on all interfaces.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the root guard feature, use the **no spanning-tree guard** or the **spanning-tree guard none** interface configuration command.

Enabling EtherChannel Guard

Use the EtherChannel guard feature to detect a misconfigured EtherChannel when Catalyst 3550 switch interfaces are configured as an EtherChannel while interfaces on the remote device are not, or not all the interfaces on the remote device are in the same EtherChannel.

Beginning in privileged EXEC mode, follow these steps to enable EtherChannel guard:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**spanning-tree etherchannel guard misconfig**	Enable the EtherChannel guard feature, and display an error message when a loop caused by a channel misconfiguration is detected. This feature is enabled by default.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show spanning-tree summary**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the EtherChannel guard feature, use the **no spanning-tree etherchannel guard misconfig** global configuration command.

To determine which local ports are involved in the misconfiguration and in the err-disabled state, enter the **show interfaces status err-disabled** privileged EXEC command. To check the EtherChannel configuration on the remote device, enter the **show etherchannel summary** privileged EXEC command on the remote device.

After you correct the configuration, enter the **shutdown** and the **no shutdown** interface configuration commands on the associated port-channel interface.

CHAPTER 11

Configuring IGMP Snooping and MVR

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on your switch, including an application of local IGMP snooping, Multicast VLAN Registration (MVR). It also includes procedures for controlling multicast group membership by using IGMP filtering.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the *Cisco IOS Release Network Protocols Command Reference, Part 1, for Release 12.1*.

This chapter consists of these sections:

- Understanding IGMP Snooping, page 11-1
- Configuring IGMP Snooping, page 11-5
- Displaying IGMP Snooping Information, page 11-9
- Understanding Multicast VLAN Registration, page 11-12
- Configuring MVR, page 11-14
- Displaying MVR Information, page 11-18
- Configuring IGMP Filtering, page 11-20
- Displaying IGMP Filtering Configuration, page 11-24

Note For MAC addresses that map to IP multicast groups, you can either manage them through features such as IGMP snooping and MVR, or you can use static MAC addresses. However, you cannot use both methods simultaneously. Therefore, before using IGMP snooping or MVR, you should remove all statically configured MAC addresses that map to IP multicast groups.

Understanding IGMP Snooping

Layer 2 switches can use IGMP snooping to constrain the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. As the name implies, IGMP snooping requires the LAN switch to snoop on the IGMP transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group,

the switch adds the host port number to the forwarding table entry; when it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

Note For more information on IP multicast and IGMP, refer to RFC 1112 and RFC 2236.

The multicast router (which could be a Catalyst 3550 switch with the enhanced multilayer software image) sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch forwards only one join request per IP multicast group to the multicast router. It creates one entry per VLAN in the Layer 2 forwarding table for each MAC group from which it receives an IGMP join request.

Layer 2 multicast groups learned through IGMP snooping are dynamic. However, you can statically configure MAC multicast groups by using the **ip igmp snooping vlan static** global configuration command. If you specify group membership for a multicast group address statically, your setting supersedes any automatic manipulation by IGMP snooping. Multicast group membership lists can consist of both user-defined and IGMP snooping-learned settings.

Note If a spanning-tree VLAN topology change occurs, the IGMP snooping-learned multicast groups on the VLAN are purged. Enabling spanning-tree Port Fast on direct-to-desktop ports stops STP topology change notifications from being generated.

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from the router, it forwards the query to all ports in the VLAN. Hosts wanting to join the multicast group respond by sending a join message to the switch. The switch CPU creates a multicast forwarding-table entry for the group if it is not already present. The CPU also adds the interface where the join message was received to the forwarding-table entry. The host associated with that interface receives multicast traffic for that multicast group. See Figure 11-1.

Figure 11-1 Initial IGMP Join Message

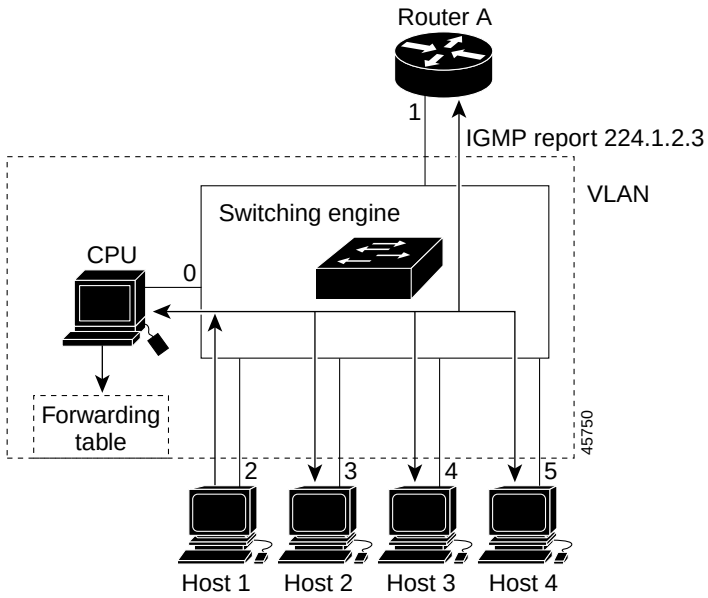

Router A sends a general query to the switch, which forwards the query to ports 2 through 5, all members of the same VLAN. Host 1 wants to join multicast group 224.1.2.3 and multicasts an IGMP membership report (IGMP join message) to the group with the equivalent MAC destination address of 0x0100.5E01.0203. When the CPU receives the IGMP report multicast by Host 1, the CPU uses the information in the IGMP report to set up a forwarding-table entry, as shown in Table 11-1, that includes the port numbers of Host 1, the router, and the switch internal CPU.

Table 11-1 IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
0100.5exx.xxxx	IGMP	0
0100.5e01.0203	!IGMP	1, 2

Note that the switch hardware can distinguish IGMP information packets from other packets for the multicast group.

- The first entry in the table tells the switching engine to send IGMP packets to only the switch CPU. This prevents the CPU from becoming overloaded with multicast frames.
- The second entry tells the switching engine to send frames addressed to the 0x0100.5E01.0203 multicast MAC address that are not IGMP packets (!IGMP) to the router and to the host that has joined the group.

If another host (for example, Host 4) sends an unsolicited IGMP join message for the same group (Figure 11-2), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in Table 11-2. Note that because the forwarding table directs IGMP messages to only the CPU, the message is not flooded to other ports on the switch. Any multicast traffic is forwarded to the group and not to the CPU.

Figure 11-2 Second Host Joining a Multicast Group

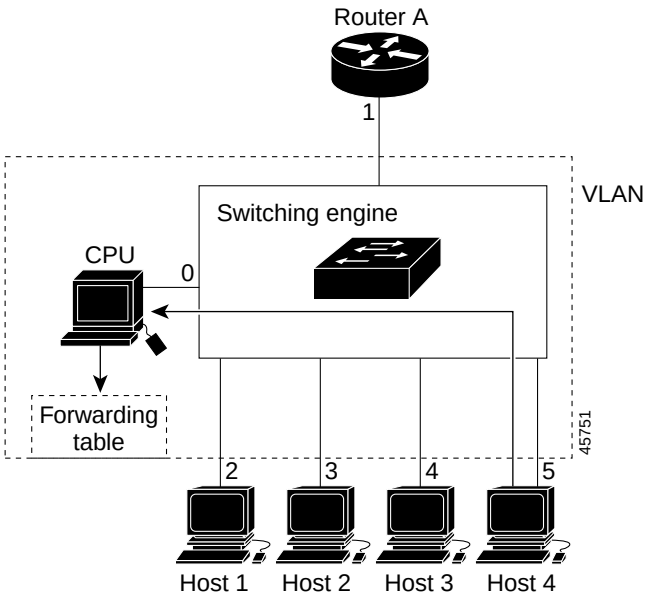

Table 11-2 Updated IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
0100.5exx.xxxx	IGMP	0
0100.5e01.0203	!IGMP	1, 2, 5

Leaving a Multicast Group

The router sends periodic multicast general queries and the switch forwards these queries through all ports in the VLAN. Interested hosts respond to the queries. If at least one host in the VLAN wishes to receive multicast traffic, the router continues forwarding the multicast traffic to the VLAN. The switch forwards multicast group traffic to only those hosts listed in the forwarding table for that Layer 2 multicast group.

When hosts want to leave a multicast group, they can either silently leave, or they can send a leave message. When the switch receives a leave message from a host, it sends out a MAC-based general query to determine if any other devices connected to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

Immediate-Leave Processing

The switch uses IGMP snooping Immediate-Leave processing to remove from the forwarding table an interface that sends a leave message without the switch sending MAC-based general queries to the interface. The VLAN interface is pruned from the multicast tree for the multicast group specified in the original leave message. Immediate-Leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are simultaneously in use.

Note You should only use the Immediate-Leave processing feature on VLANs where a single host is connected to each port. If Immediate Leave is enabled in VLANs where more than one host is connected to a port, some hosts might be inadvertently dropped. Immediate Leave is supported with only IGMP version 2 hosts.

Configuring IGMP Snooping

IGMP snooping allows switches to examine IGMP packets and make forwarding decisions based on their content. To enable IGMP snooping on the switch to discover external multicast routers, the Layer 3 interfaces on the routers in the VLAN must already have been configured for multicast routing. For more information, see Chapter 24, "Configuring IP Multicast Routing."

These sections describe how to configure IGMP snooping:

- Default IGMP Snooping Configuration, page 11-5
- Enabling or Disabling IGMP Snooping, page 11-5
- Setting the Snooping Method, page 11-6
- Configuring a Multicast Router Port, page 11-7
- Configuring a Host Statically to Join a Group, page 11-8
- Enabling IGMP Immediate-Leave Processing, page 11-9

Default IGMP Snooping Configuration

Table 11-3 shows the default IGMP snooping configuration.

Table 11-3 Default IGMP Snooping Configuration

Feature	Default Setting
IGMP snooping	Enabled globally and per VLAN
Multicast routers	None configured
Multicast router learning (snooping) method	PIM-DVMRP
IGMP snooping Immediate Leave	Disabled
Static groups	None configured

Enabling or Disabling IGMP Snooping

By default, IGMP snooping is globally enabled on the switch. When globally enabled or disabled, it is also enabled or disabled in all existing VLAN interfaces. IGMP snooping is by default enabled on all VLANs, but can be enabled and disabled on a per-VLAN basis. After you configure the VLAN interface for multicast routing, no configuration is needed for the switch to dynamically access external multicast routers by using IGMP snooping.

Global IGMP snooping overrides the VLAN IGMP snooping. If global snooping is disabled, you cannot enable VLAN snooping. If global snooping is enabled, you can enable or disable VLAN snooping.

Beginning in privileged EXEC mode, follow these steps to globally enable IGMP snooping on the switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip igmp snooping**	Globally enable IGMP snooping in all existing VLAN interfaces.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To globally disable IGMP snooping on all VLAN interfaces, use the **no ip igmp snooping** global configuration command.

Beginning in privileged EXEC mode, follow these steps to enable IGMP snooping on a VLAN interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip igmp snooping vlan** *vlan-id*	Enable IGMP snooping on the VLAN interface.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable IGMP snooping on a VLAN interface, use the **no ip igmp snooping vlan** *vlan-id* global configuration command for the specified VLAN number.

Setting the Snooping Method

Multicast-capable router ports are added to the forwarding table for every Layer 2 multicast entry. The switch learns of such ports through one of these methods:

- Snooping on IGMP queries, Protocol Independent Multicast (PIM) packets, and Distance Vector Multicast Routing Protocol (DVMRP) packets
- Listening to Cisco Group Management Protocol (CGMP) packets from other routers
- Statically connecting to a multicast router port with the **ip igmp snooping mrouter** global configuration command

You can configure the switch either to snoop on IGMP queries and PIM/DVMRP packets or to listen to CGMP self-join or proxy-join packets. By default, the switch snoops on PIM/DVMRP packets on all VLANs. To learn of multicast router ports through only CGMP packets, use the **ip igmp snooping vlan** *vlan-id* **mrouter learn cgmp** global configuration command. When this command is issued, the router listens to only CGMP self-join and CGMP proxy-join packets and no other CGMP packets. To learn of multicast router ports through only PIM-DVMRP packets, use the **ip igmp snooping vlan** *vlan-id* **mrouter learn pim-dvmrp** global configuration command.

Note If you want to use CGMP as the learning method and no multicast routers in the VLAN are CGMP proxy-enabled, you must enter the **ip cgmp router-only** command to dynamically access the router. For more information, see Chapter 24, "Configuring IP Multicast Routing."

Beginning in privileged EXEC mode, follow these steps to alter the method in which a VLAN interface dynamically accesses a multicast router:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip igmp snooping vlan** *vlan-id* **mrouter learn** {**cgmp** | **pim-dvmrp**}	Enable IGMP snooping on a VLAN. The VLAN ID range is 1 to 1001. Specify the multicast router learning method: • **cgmp**—Listen for CGMP packets. This method is useful for reducing control traffic. • **pim-dvmrp**—Snoop on IGMP queries and PIM-DVMRP packets. This is the default.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip igmp snooping**	Verify the configuration.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

This example shows how to configure IGMP snooping to use CGMP packets as the learning method:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 mrouter learn cgmp
Switch(config)# end
Switch# show ip igmp snooping vlan 1
vlan 1
----------
 IGMP snooping is globally enabled
 IGMP snooping is enabled on this Vlan
 IGMP snooping immediate-leave is disabled on this Vlan
 IGMP snooping mrouter learn mode is cgmp on this Vlan
 IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

To return to the default learning method, use the **no ip igmp snooping vlan** *vlan-id* **mrouter learn cgmp** global configuration command.

Configuring a Multicast Router Port

To add a multicast router port (add a static connection to a multicast router), use the **ip igmp snooping vlan mrouter** global configuration command on the switch.

Note Static connections to multicast routers are supported only on switch ports.

Beginning in privileged EXEC mode, follow these steps to enable a static connection to a multicast router:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip igmp snooping vlan** *vlan-id* **mrouter interface** *interface-id*	Specify the multicast router VLAN ID (1 to 1001), and specify the interface to the multicast router.
Step 3	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 4	**show ip igmp snooping mrouter** [**vlan** *vlan-id*]	Verify that IGMP snooping is enabled on the VLAN interface.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a multicast router port from the VLAN, use the **no ip igmp snooping vlan** *vlan-id* **mrouter interface** *interface-id* global configuration command.

This example shows how to enable a static connection to a multicast router and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 200 mrouter interface gigabitethernet0/2
Switch(config)# end
Switch# show ip igmp snooping mrouter vlan 200
vlan            ports
-----+---------------------------------------
 200           Gi0/2(static)
```

Configuring a Host Statically to Join a Group

Hosts or Layer 2 ports normally join multicast groups dynamically, but you can also statically configure a host on an interface.

Beginning in privileged EXEC mode, follow these steps to add a Layer 2 port as a member of a multicast group:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**ip igmp snooping vlan** *vlan-id* **static** *mac-address* **interface** *interface-id*	Statically configure a Layer 2 port as a member of a multicast group: • *vlan-id* is the multicast group VLAN ID. • *mac-address* is the group MAC address. • *interface-id* is the member port.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip igmp snooping mrouter vlan** *vlan-id* or **show mac-address-table multicast vlan** *vlan-id*	Verify that the member port is a member of the VLAN multicast group. Verify the member port and the MAC address
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the Layer 2 port from the multicast group, use the **no ip igmp snooping vlan** *vlan-id* **static** *mac-address* **interface** *interface-id* global configuration command.

This example shows how to statically configure a host on an interface and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 static 0100.5e00.0203 interface
gigabitethernet0/11
Switch(config)# end
Switch# show mac-address-table multicast vlan 1
Vlan    Mac Address       Type       Ports
----    -----------       ----       -----
   1    0100.5e00.0203    USER       Gi0/11
```

Enabling IGMP Immediate-Leave Processing

When you enable IGMP Immediate-Leave processing, the switch immediately removes a port when it detects an IGMP version 2 leave message on that port. You should use the Immediate-Leave feature only when there is a single receiver present on every port in the VLAN.

Immediate Leave is supported with only IGMP version 2 hosts.

Beginning in privileged EXEC mode, follow these steps to enable IGMP Immediate-Leave processing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**ip igmp snooping vlan** *vlan-id* **immediate-leave**	Enable IGMP Immediate-Leave processing on the VLAN interface.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip igmp snooping vlan** *vlan-id*	Verify that Immediate Leave is enabled on the VLAN.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable IGMP Immediate-Leave on a VLAN, use the **no ip igmp snooping vlan** *vlan-id* **immediate-leave** global configuration command.

This example shows how to enable IGMP immediate-leave processing on VLAN 130 and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
Switch# show ip igmp snooping vlan 130
vlan 130
----------
 IGMP snooping is globally enabled
 IGMP snooping is enabled on this Vlan
 IGMP snooping immediate-leave is enabled on this Vlan
 IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
 IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

Displaying IGMP Snooping Information

You can display IGMP snooping information for dynamically learned and statically configured router ports and VLAN interfaces. You can also display MAC address multicast entries for a VLAN configured for IGMP snooping.

To display IGMP snooping information, use one or more of the privileged EXEC commands in Table 11-4.

Table 11-4 Commands for Displaying IGMP Snooping Information

Command	Purpose
show ip igmp snooping [**vlan** *vlan-id*]	Display the snooping configuration information for all VLANs on the switch or for a specified VLAN. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN.
show ip igmp snooping mrouter [**vlan** *vlan-id*]	Display information on dynamically learned and manually configured multicast router interfaces. **Note** When you enable IGMP snooping, the switch automatically learns the interface to which a multicast router is connected. These are dynamically learned interfaces. (Optional) Enter **vlan** *vlan-id* to display information for a single VLAN.
show mac-address-table multicast [**vlan** *vlan-id*] [**user** | **igmp-snooping**] [**count**]	Display the Layer 2 MAC address table entries for a VLAN. The keywords are all optional and limit the display as shown: • **vlan** *vlan-id*—Displays only the specified multicast group VLAN. • **user**—Displays only the user-configured multicast entries. • **igmp-snooping**—Displays only entries learned through IGMP snooping. • **count**—Displays only the total number of entries for the selected criteria, not the actual entries.

This is an example of output from the **show ip igmp snooping** privileged EXEC command for all VLAN interfaces on the switch:

```
Switch# show ip igmp snooping
vlan 1
----------
 IGMP snooping is globally enabled
 IGMP snooping is enabled on this Vlan
 IGMP snooping immediate-leave is disabled on this Vlan
 IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
 IGMP snooping is running in IGMP_ONLY mode on this Vlan
vlan 2
----------
 IGMP snooping is globally enabled
 IGMP snooping is enabled on this Vlan
 IGMP snooping immediate-leave is disabled on this Vlan
 IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
 IGMP snooping is running in IGMP_ONLY mode on this Vlan
vlan 10
----------
 IGMP snooping is globally enabled
 IGMP snooping is enabled on this Vlan
 IGMP snooping immediate-leave is disabled on this Vlan
 IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
 IGMP snooping is running in IGMP_ONLY mode on this Vlan
```

This is an example of output from the **show ip igmp snooping** privileged EXEC command for a specific VLAN interface:

```
Switch# show ip igmp snooping vlan 1
vlan 1
----------
  IGMP snooping is globally enabled
  IGMP snooping is disabled on this Vlan
  IGMP snooping immediate-leave is disabled on this Vlan
  IGMP snooping mrouter learn mode is pim-dvmrp on this Vlan
```

This is an example of output from the **show ip igmp snooping mrouter** privileged EXEC command:

```
Switch# show ip igmp snooping mrouter
vlan            ports
-----+---------------------------------------
  1          Gi0/1,Gi0/2,Router
  2          Gi0/3,Gi0/4
```

This is an example of output from the **show ip igmp snooping mrouter** privileged EXEC command for VLAN 1:

```
Switch# show ip igmp snooping mrouter vlan 1
Vlan    ports
----    -----
   1    Gi0/1(dynamic)
   1    Gi0/2(dynamic)
```

This example shows how to display the Layer 2 multicast entries for VLAN 1:

```
Switch# show mac-address-table multicast vlan 1
vlan   mac address     type         ports
-----+---------------+--------+---------+---+--------------------------------
  1  0100.5e02.0203    user           Gi0/1,Gi0/2
  1  0100.5e00.0127    igmp           Gi0/1,Gi0/2
  1  0100.5e00.0128    user           Gi0/1,Gi0/2
  1  0100.5e00.0001    igmp           Gi0/1,Gi0/2
```

This is an example of output from the **show mac-address-table multicast count** privileged EXEC command for the switch:

```
Switch# show mac-address-table multicast count

Multicast MAC Entries for all vlans:    10
```

This is an example of output from the **show mac-address-table multicast count** privileged EXEC command for a VLAN:

```
Switch# show mac-address-table multicast vlan 1 count

Multicast MAC Entries for vlan 1:
```

This example shows how to display only the user-configured multicast entries for VLAN 1:

```
Switch# show mac-address-table multicast vlan 1 user
vlan   mac address     type         ports
-----+---------------+--------+---------+---+--------------------------------
  1  0100.5e02.0203    user           Gi0/1,Gi0/2
  1  0100.5e00.0128    user           Gi0/1,Gi0/2
```

This example shows how to display the total number of entries learned by IGMP snooping for VLAN 1:

```
Switch# show mac-address-table multicast vlan 1 igmp-snooping count

Number of user programmed entries:    2
```

Understanding Multicast VLAN Registration

Multicast VLAN Registration (MVR) is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service provider network (for example, the broadcast of multiple television channels over a service-provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast VLAN. It allows the single multicast VLAN to be shared in the network while subscribers remain in separate VLANs. MVR provides the ability to continuously send multicast streams in the multicast VLAN, but to isolate the streams from the subscriber VLANs for bandwidth and security reasons.

MVR assumes that subscriber ports subscribe and unsubscribe (join and leave) these multicast streams by sending out IGMP join and leave messages. These messages can originate from an IGMP version-2-compatible host with an Ethernet connection. Although MVR operates on the underlying mechanism of IGMP snooping, the two features operate independently of each other. One can be enabled or disabled without affecting the behavior of the other feature. However, if IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping.

The switch CPU identifies the MVR IP multicast streams and their associated MAC addresses in the switch forwarding table, intercepts the IGMP messages, and modifies the forwarding table to include or remove the subscriber as a receiver of the multicast stream, even though the receivers might be in a different VLAN from the source. This forwarding behavior selectively allows traffic to cross between different VLANs.

The Catalyst 3550 switch has these modes of MVR operation: dynamic and compatible.

- When operating in MVR dynamic mode, the switch performs standard IGMP snooping. IGMP information packets are sent to the switch CPU, but multicast data packets are not sent to the CPU. Dynamic mode allows the multicast router to run normally because the switch sends the IGMP join messages to the router, and the router forwards multicast streams for a particular group to an interface only if it has received a join message from the interface for the group. Receiver ports are treated as members of the multicast VLAN for MVR multicast control and data traffic. IGMP reports for MVR groups are sent out source ports in the multicast VLAN.
- When in MVR compatible mode, MVR on the Catalyst 3550 switch interoperates with MVR on Catalyst 3500 XL and Catalyst 2900 XL switches. It works the same as dynamic mode for all multicast data packets and IGMP query and leave packets. However, received IGMP report packets for MVR groups are not sent out on the multicast VLAN source ports. In contrast to dynamic mode, the switch does not send join messages to the router. The router must be statically configured for the interface to receive the multicast stream. Therefore, in this mode, MVR does not support dynamic membership joins on source ports.

Using MVR in a Multicast Television Application

In a multicast television application, a PC or a television with a set-top box can receive the multicast stream. Multiple set-top boxes or PCs can be connected to one subscriber port, which is a switch port configured as an MVR receiver port. See Figure 11-3. DHCP assigns an IP address to the set-top box or the PC. When a subscriber selects a channel, the set-top box or PC sends an IGMP report to the S1 switch to join the appropriate multicast. If the IGMP report matches one of the configured multicast MAC addresses, the switch CPU modifies the hardware address table to include this receiver port and VLAN as a forwarding destination of the specified multicast stream when it is received from the multicast VLAN. Uplink ports that send and receive multicast data to and from the multicast VLAN are called MVR source ports.

When a subscriber changes channels or turns off the television, the set-top box sends an IGMP leave message for the multicast stream. The switch CPU sends an IGMP group-specific query through the receiver port VLAN. If there is another set-top box in the VLAN still subscribing to this group, that set-top box must respond within the maximum response time. If the CPU does not receive a response, it eliminates the receiver port as a forwarding destination for this group.

If the Immediate-Leave feature is enabled on a receiver port, the port leaves a multicast group more quickly. Without Immediate Leave, when the switch receives an IGMP leave message from a subscriber on a receiver port, it sends out an IGMP query on that port and waits for IGMP group membership reports. If no reports are received in a configured time period, the receiver port is removed from multicast group membership. With Immediate Leave, an IGMP query is not sent from the receiver port on which the IGMP leave was received. As soon as the leave message is received, the receiver port is removed from multicast group membership, which speeds up leave latency. Enable the Immediate Leave feature only on receiver ports to which a single receiver device is connected.

Figure 11-3 Multicast VLAN Registration Example

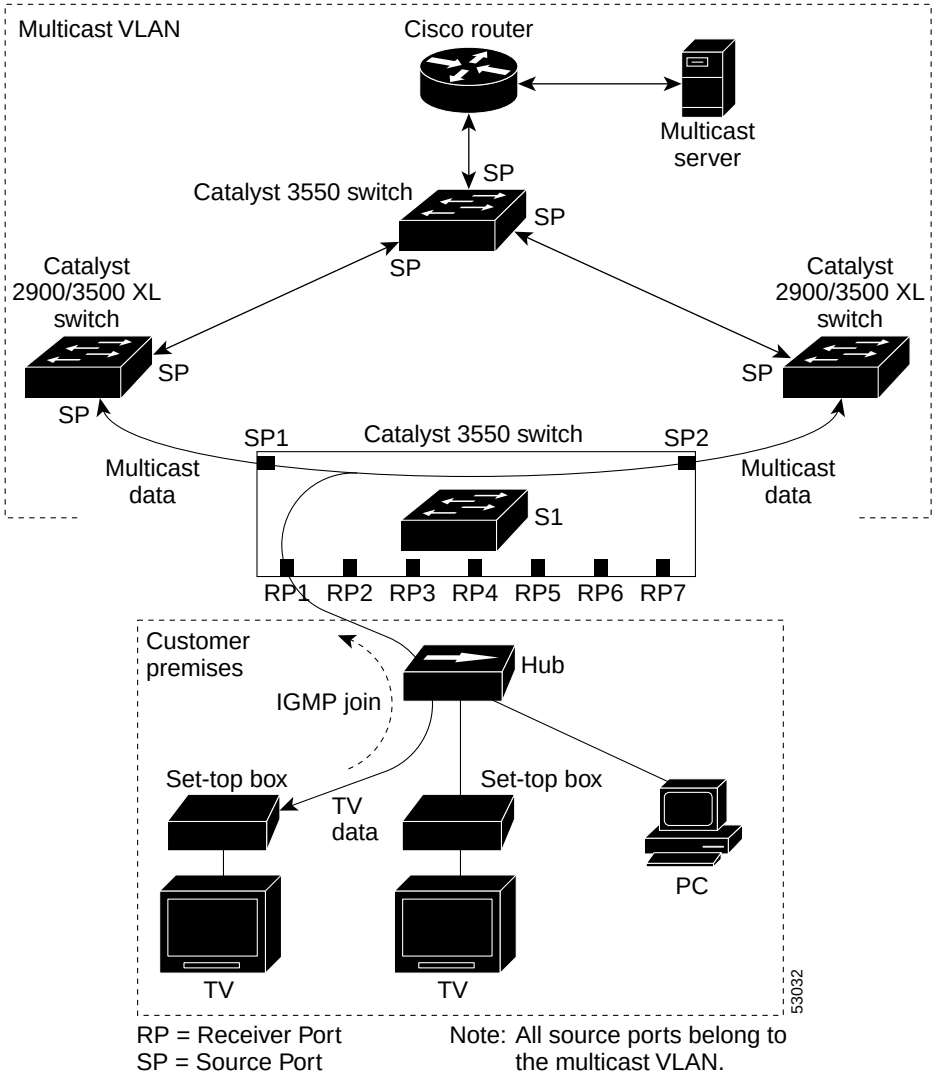

MVR eliminates the need to duplicate television-channel multicast traffic for subscribers in each VLAN. Multicast traffic for all channels is only sent around the VLAN trunk once—only on the multicast VLAN. Although the IGMP leave and join message in the VLAN to which the subscriber port is assigned. These messages dynamically register for streams of multicast traffic in the multicast VLAN on the Layer 3 device. The access layer switch (S1 switch) modifies the forwarding behavior to allow the traffic to be forwarded from the multicast VLAN to the subscriber port in a different VLAN, selectively allowing traffic to cross between two VLANs.

IGMP reports are sent to the same MAC addresses as the multicast data. The S1 CPU must capture all IGMP join and leave messages from receiver ports and forward them to the multicast VLAN of the source (uplink) port.

Configuring MVR

These sections include basic MVR configuration information:

- Configuration Guidelines and Limitations, page 11-14
- Default MVR Configuration, page 11-15
- Configuring MVR Global Parameters, page 11-15
- Configuring MVR Interfaces, page 11-16

Configuration Guidelines and Limitations

Follow these guidelines when configuring MVR:

- Receiver ports cannot be trunk ports. Receiver ports on a switch can be in different VLANs, but should not belong to the multicast VLAN.
- The maximum number of multicast entries that can be configured on a switch (that is, the maximum number of television channels that can be received) is 256.
- Each channel is one multicast stream destined for a unique IP multicast address. These IP addresses cannot alias between themselves or with the reserved IP multicast addresses (in the range 224.0.0.xxx).
- Multicast routing and MVR cannot coexist on a switch. If you enable multicast routing and a multicast routing protocol while MVR is enabled, MVR is disabled, and you receive a warning message. If you try to enable MVR while multicast routing and a multicast routing protocol are enabled, the operation to enable MVR is cancelled, and you receive an error message.

Note For complete syntax and usage information for the commands used in this section, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Default MVR Configuration

Table 11-5 shows the default MVR configuration.

Table 11-5 Default MVR Configuration

Feature	Default Setting
MVR	Disabled globally and per interface
Multicast addresses	None configured
Query response time	0.5 second
Multicast VLAN	VLAN 1
Mode	Compatible
Interface (per port) default	Neither a receiver nor a source port
Immediate Leave	Disabled on all ports

Configuring MVR Global Parameters

You do not need to set the optional MVR parameters if you choose to use the default settings. If you do want to change the default parameters (except for the MVR VLAN), you must first enable MVR.

Beginning in privileged EXEC mode, follow these steps to configure MVR parameters:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mvr**	Enable MVR on the switch.
Step 3	**mvr group** *ip-address* [*count*]	Configure an IP multicast address on the switch or use the *count* parameter to configure a contiguous series of MVR group addresses (the range for *count* is 1 to 256; the default is 1). Any multicast data sent to this address is sent to all source ports on the switch and all receiver ports that have elected to receive data on that multicast address. Each multicast address would correspond to one television channel. **Note** Each IP address translates to a multicast 48-bit MAC address. If an IP address being configured translates (aliases) to a previously configured MAC address or to any reserved multicast MAC addresses, the command fails.
Step 4	**mvr querytime** *value*	(Optional) Define the maximum time to wait for IGMP report memberships on a receiver port before removing the port from multicast group membership. The value is in units of tenths of a second. The range is from 1 to 100 and the default is 5 tenths or one-half second.
Step 5	**mvr vlan** *vlan-id*	(Optional) Specify the VLAN in which multicast data is received; all source ports must belong to this VLAN. The VLAN range is 1 to 1001. The default is VLAN 1.

	Command	Purpose
Step 6	**mvr mode** {**dynamic** | **compatible**}	(Optional) Specify the MVR mode of operation: • **dynamic**—Allows dynamic MVR membership on source ports. • **compatible**—Is compatible with Catalyst 3500 XL and Catalyst 2900 XL switches and does not support IGMP dynamic joins on source ports. The default is **compatible** mode.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show mvr** or **show mvr members**	Verify the configuration.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the switch to its default settings, use the **no mvr** [**mode** | **group** *ip-address* | **querytime** | **vlan**] global configuration commands.

This example shows how to enable MVR, configure the MVR group address, set the query time to 1 second (10 tenths), specify the MVR multicast VLAN as VLAN 22, set the MVR mode as dynamic, and verify the results:

```
Switch(config)# mvr
Switch(config)# mvr group 228.1.23.4
Switch(config)# mvr querytime 10
Switch(config)# mvr vlan 22
Switch(config)# mvr mode dynamic
Switch(config)# end
Switch# show mvr
MVR Running: TRUE
MVR multicast vlan: 22
MVR Max Multicast Groups: 256
MVR Current multicast groups: 1
MVR Global query response time: 10 (tenths of sec)
MVR Mode: dynamic
```

You can use the **show mvr members** privileged EXEC command to verify the MVR multicast group addresses on the switch.

Configuring MVR Interfaces

Beginning in privileged EXEC mode, follow these steps to configure MVR interfaces:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mvr**	Enable MVR on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and enter the type and number of the port to configure, for example, **gi0/1** or **gigabitethernet 0/1** for Gigabit Ethernet port 1.

	Command	Purpose
Step 4	**mvr type** {**source** | **receiver**}	Configure an MVR port as one of these: • **source**—Configure uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. All source ports on a switch belong to the single multicast VLAN. • **receiver**—Configure a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group, either statically or by using IGMP leave and join messages. Receiver ports cannot belong to the multicast VLAN. The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.
Step 5	**mvr vlan** *vlan-id* **group** *ip-address*	(Optional) Statically configure a port to receive multicast traffic sent to the multicast VLAN and the IP multicast address. A port statically configured as a member of a group remains a member of the group until statically removed. **Note** In compatible mode, this command applies to only receiver ports. In dynamic mode, it applies to receiver ports and source ports. Receiver ports can also dynamically join multicast groups by using IGMP join and leave messages.
Step 6	**mvr immediate**	(Optional) Enable the Immediate Leave feature of MVR on the port. **Note** This command applies to only receiver ports and should only be enabled on receiver ports to which a single receiver device is connected.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show mvr** **show mvr interface** or **show mvr members**	Verify the configuration.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the interface to its default settings, use the **no mvr** [**type** | **immediate** | **vlan** *vlan-id* | **group**] interface configuration commands.

This example shows how to configure Gigabit Ethernet port 0/1 as a receiver port, statically configure the port to receive multicast traffic sent to the multicast group address, configure Immediate Leave on the interface, and verify the results.

```
Switch(config)# mvr
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# mvr type receiver
Switch(config-if)# mvr vlan 22 group 228.1.23.4
Switch(config-if)# mvr immediate
Switch(config)# end
Switch# show mvr interface gigabitethernet0/2
Type: RECEIVER Status: ACTIVE Immediate Leave: ENABLED
```

This is an example of output from the **show mvr interface** privileged EXEC command when the **member** keyword is included:

```
Switch# show mvr interface gigabitethernet0/6 member
239.255.0.0     DYNAMIC ACTIVE
239.255.0.1     DYNAMIC ACTIVE
239.255.0.2     DYNAMIC ACTIVE
239.255.0.3     DYNAMIC ACTIVE
239.255.0.4     DYNAMIC ACTIVE
239.255.0.5     DYNAMIC ACTIVE
239.255.0.6     DYNAMIC ACTIVE
239.255.0.7     DYNAMIC ACTIVE
239.255.0.8     DYNAMIC ACTIVE
239.255.0.9     DYNAMIC ACTIVE
```

Displaying MVR Information

You can display MVR information for the switch or for a specified interface.

Beginning in privileged EXEC mode, use the commands in Table 11-6 to display MVR configuration:

Table 11-6 Commands for Displaying MVR Information

show mvr	Displays MVR status and values for the switch—whether MVR is enabled or disabled, the multicast VLAN, the maximum (256) and current (0 through 256) number of multicast groups, the query response time, and the MVR mode.
show mvr interface [*interface-id*] [**members** [**vlan** *vlan-id*]]	Displays all MVR interfaces and their MVR configurations. When a specific interface is entered, displays this information: • Type—Receiver or Source • Status—One of these: – Active means the port is part of a VLAN. – Up/Down means that the port is forwarding or nonforwarding. – Inactive means that the port is not part of any VLAN. • Immediate Leave—Enabled or Disabled If the **members** keyword is entered, displays all multicast group members on this port or, if a VLAN identification is entered, all multicast group members on the VLAN. The VLAN ID can be from 1 to 1001.
show mvr members [*ip-address*]	Displays all receiver ports that are members of any IP multicast group or the specified IP multicast group IP address.

This is an example of output from the **show mvr** privileged EXEC command:

```
Switch# show mvr
MVR Running: TRUE
MVR multicast vlan: 1
MVR Max Multicast Groups: 256
MVR Current multicast groups: 256
MVR Global query response time: 5 (tenths of sec)
MVR Mode: compatible
```

This is an example of output from the **show mvr interface** privileged EXEC command:

```
Switch# show mvr interface
Port    Type            Status          Immediate Leave
----    ----            -------         ---------------
Gi0/1   SOURCE          ACTIVE/UP       DISABLED
Gi0/2   SOURCE          ACTIVE/UP       DISABLED
Gi0/3   RECEIVER        ACTIVE/UP       DISABLED
Gi0/4   RECEIVER        ACTIVE/UP       DISABLED
Gi0/5   RECEIVER        ACTIVE/UP       ENABLED
Gi0/6   RECEIVER        ACTIVE/UP       DISABLED
Gi0/7   RECEIVER        ACTIVE/UP       ENABLED
Gi0/8   RECEIVER        ACTIVE/UP       DISABLED
```

This is an example of output from the **show mvr interface** privileged EXEC command for a specified interface:

```
Switch# show mvr interface gigabitethernet0/2
Type: RECEIVER Status: ACTIVE Immediate Leave: DISABLED
```

This is an example of output from the **show mvr interface** privileged EXEC command when the **members** keyword is included:

```
Switch# show mvr interface gigabitethernet0/6 members
239.255.0.0     DYNAMIC ACTIVE
239.255.0.1     DYNAMIC ACTIVE
239.255.0.2     DYNAMIC ACTIVE
239.255.0.3     DYNAMIC ACTIVE
239.255.0.4     DYNAMIC ACTIVE
239.255.0.5     DYNAMIC ACTIVE
239.255.0.6     DYNAMIC ACTIVE
239.255.0.7     DYNAMIC ACTIVE
239.255.0.8     DYNAMIC ACTIVE
239.255.0.9     DYNAMIC ACTIVE
```

This is an example of output from the **show mvr members** privileged EXEC command:

```
Switch# show mvr members
MVR Group IP    Status          Members
------------    ------          -------
239.255.0.1     ACTIVE          Gi0/1(d), Gi0/5(s)
239.255.0.2     INACTIVE        None
239.255.0.3     INACTIVE        None
239.255.0.4     INACTIVE        None
239.255.0.5     INACTIVE        None
239.255.0.6     INACTIVE        None
239.255.0.7     INACTIVE        None
239.255.0.8     INACTIVE        None
239.255.0.9     INACTIVE        None
239.255.0.10    INACTIVE        None

<output truncated>

239.255.0.255   INACTIVE        None
239.255.1.0     INACTIVE        None
```

Configuring IGMP Filtering

In some environments, for example metropolitan or multiple-dwelling unit (MDU) installations, an administrator might want to control the set of multicast groups to which a user on a switch port can belong. This allows the administrator to control the distribution of multicast services, such as IP/TV, based on some type of subscription or service plan. With the IGMP filtering feature, you can filter multicast joins on a per-port basis by configuring IP multicast profiles and associating them with individual switch ports. An IGMP profile can contain one or more multicast groups and specifies whether access to the group is permitted or denied. If an IGMP profile denying access to a multicast group is applied to a switch port, the IGMP join report requesting the stream of IP multicast traffic is dropped, and the port is not allowed to receive IP multicast traffic from that group. If the filtering action permits access to the multicast group, the IGMP report from the port is forwarded for normal processing.

IGMP filtering controls only IGMP membership join reports and has no relationship with the function that directs the forwarding of IP multicast traffic. The filtering feature operates in the same manner whether CGMP or MVR is used to forward the multicast traffic.

You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

Default IGMP Filtering Configuration

Table 11-5 shows the default IGMP filtering configuration.

Table 11-7 Default MVR Configuration

Feature	Default Setting
IGMP filters	None applied
IGMP Maximum number of IGMP groups	No maximum set
IGMP profiles	None defined
IGMP profile action	Deny the range addresses

Configuring IGMP Profiles

To configure an IGMP profile, use the **ip igmp profile** global configuration command with a profile number to create an IGMP profile andto enter IGMP profile configuration mode. From this mode, you can specify the parameters of the IGMP profile to be used for filtering IGMP join requests from a port. When you are in IGMP profile configuration mode, you can create the profile by using these commands:

- **deny**: Specifies that matching addresses are denied; this is the default condition.
- **exit**: Exits from igmp-profile configuration mode.
- **no**: Negates a command or sets its defaults.
- **permit**: Specifies that matching addresses are permitted.
- **range**: Specifies a range of IP addresses for the profile. You can enter a single IP address or a range with a start and an end address.

The default is for the switch to have no IGMP profiles configured. When a profile is configured, if neither the **permit** nor **deny** keyword is included, the default is to deny access to the range of IP addresses.

Beginning in privileged EXEC mode, follow these steps to create an IGMP profile:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip igmp profile** *profile number*	Enter IGMP profile configuration mode, and assign a number to the profile you are configuring. The range is from 1 to 4294967295.
Step 3	**permit** | **deny**	(Optional) Set the action to permit or deny access to the IP multicast address. If no action is configured, the default for the profile is to deny access.
Step 4	**range** *ip multicast address*	Enter the IP multicast address or range of IP multicast addresses to which access is being controlled. If entering a range, enter the low IP multicast address, a space, and the high IP multicast address. You can use the **range** command multiple times to enter multiple addresses or ranges of addresses.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show ip igmp profile** *profile number*	Verify the profile configuration.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete a profile, use the **no ip igmp profile** *profile number* global configuration command.

To delete an IP multicast address or range of IP multicast addresses, use the **no range** *ip multicast address* IGMP profile configuration command.

This example shows how to create IGMP profile 4 allowing access to the single IP multicast address and how to verify the configuration. If the action was to deny (the default), it would not appear in the **show ip igmp profile** output display.

```
Switch # config t
Switch(config) # ip igmp profile 4
Switch(config-igmp-profile)# permit
Switch(config-igmp-profile)# range 229.9.9.0
Switch(config-igmp-profile)# end
Switch# show ip igmp profile 4
IGMP Profile 4
    permit
    range 229.9.9.0 229.9.9.0
```

Applying IGMP Profiles

To control access as defined in an IGMP profile, use the **ip igmp filter** interface configuration command to apply the profile to the appropriate interfaces. You can apply IGMP profiles to Layer 2 ports only; you cannot apply IGMP profiles to routed ports or SVIs. You can apply a profile to multiple interfaces, but each interface can only have one profile applied to it.

Beginning in privileged EXEC mode, follow these steps to apply an IGMP profile to a switch port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the physical interface to configure, for example **fastethernet0/3**. The interface must be a Layer 2 port.
Step 3	**ip igmp filter** *profile number*	Apply the specified IGMP profile to the interface. The profile number can be from 1 to 4294967295.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running configuration interface** *interface-id*	Verify the configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a profile from an interface, use the **no ip igmp filter** *profile number* interface configuration command.

This example shows how to apply IGMP profile 4 to an interface and verify the configuration.

```
Switch # config t
Switch(config)# interface fastethernet0/12
Switch(config-if)# ip igmp filter 4
Switch(config-if)# end
Switch# show running-config interface fastethernet0/12
Building configuration...

Current configuration : 123 bytes
!
interface FastEthernet0/12
 no ip address
 shutdown
 snmp trap link-status
 ip igmp max-groups 25
 ip igmp filter 4
end
```

Setting the Maximum Number of IGMP Groups

You can set the maximum number of IGMP groups that a Layer 2 interface can join. Use the **no** form of this command to set the maximum back to the default, which is no limit.

This restriction can be applied to Layer 2 ports only; you cannot set a maximum number of IGMP groups on routed ports or SVIs.

Beginning in privileged EXEC mode, follow these steps to apply an IGMP profile to a switch port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the physical interface to configure, for example **gigabitethernet0/1**. The interface must be a Layer 2 port.
Step 3	**ip igmp max-groups** *number*	Set the maximum number of IGMP groups that the interface can join. The range is from 0 to 4294967294. The default is to have no maximum set.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-configuration interface** *interface-id*	Verify the configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the maximum group limitation and return to the default of no maximum, use the **no ip igmp max-groups** interface configuration command.

This example shows how to limit the number of IGMP groups that an interface can join to 25.

```
Switch# config t
Switch(config)# interface fastethernet0/12
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# end
Switch# show running-config interface fastethernet0/12
Building configuration...

Current configuration : 123 bytes
!
interface FastEthernet0/12
 no ip address
 shutdown
 snmp trap link-status
 ip igmp max-groups 25
 ip igmp filter 4
end
```

Displaying IGMP Filtering Configuration

You can display IGMP profile characteristics, and you can display the IGMP profile and maximum group configuration for all interfaces on the switch or for a specified interface.

Beginning in privileged EXEC mode, use the commands in Table 11-8 to display IGMP filtering configuration:

Table 11-8 Commands for Displaying IGMP Filtering Configuration

show ip igmp profile [*profile number*]	Displays the specified IGMP profile or all IGMP profiles defined on the switch.
show running-configuration [**interface** *interface-id*]	Displays the configuration of the specified interface or all interfaces on the switch, including (if configured) the maximum number of IGMP groups to which an interface can belong and the IGMP profile applied to the interface.

This is an example of the **show ip igmp profile** privileged EXEC command when no profile number is entered. All profiles defined on the switch are displayed.

```
Switch# show ip igmp profile
IGMP Profile 3
    range 230.9.9.0 230.9.9.0
IGMP Profile 4
    permit
    range 229.9.9.0 229.255.255.255
```

This is an example of the output from the **show running-config** privileged EXEC command when an interface is specified with IGMP maximum groups configured and IGMP profile 4 has been applied to the interface.

```
Switch# show running-config interface fastethernet0/12
Building configuration...

Current configuration : 123 bytes
!
interface FastEthernet0/12
 no ip address
 shutdown
 snmp trap link-status
 ip igmp max-groups 25
 ip igmp filter 4
end
```

Configuring Port-Based Traffic Control

This chapter describes how to configure the port-based traffic control features on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Configuring Storm Control, page 12-1
- Configuring Protected Ports, page 12-5
- Configuring Port Blocking, page 12-6
- Configuring Port Security, page 12-8
- Displaying Port-Based Traffic Control Settings, page 12-11

Configuring Storm Control

These sections include storm control configuration information and procedures:

- Understanding Storm Control, page 12-1
- Default Storm Control Configuration, page 12-3
- Enabling Storm Control, page 12-3
- Disabling Storm Control, page 12-4

Understanding Storm Control

Storm control prevents switchports on a LAN from being disrupted by a broadcast, multicast, or unicast storm on one of the physical interfaces. A LAN storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation or in the network configuration can cause a storm.

Storm control (or traffic suppression) monitors incoming traffic statistics over a time period and compares the measurement with a predefined suppression level threshold. The threshold represents the percentage of the total available bandwidth of the port. The switch supports separate storm control thresholds for broadcast, multicast, and unicast traffic. If the threshold of a traffic type is reached, further traffic of that type is suppressed until the incoming traffic falls below the threshold level.

Note When the rate of multicast traffic exceeds a set threshold, all incoming traffic (broadcast, multicast, and unicast) is dropped until the level drops below the threshold level. Only spanning-tree packets are forwarded. When broadcast and unicast thresholds are exceeded, traffic is blocked for only the type of traffic that exceeded the threshold.

When storm control is enabled, the switch monitors packets passing from an interface to the switching bus and determines if the packet is unicast, multicast, or broadcast. The switch monitors the number of broadcast, multicast, or unicast packets received within the 1-second time interval, and when a threshold for one type of traffic is reached, that type of traffic is dropped. This threshold is specified as a percentage of total available bandwidth that can be used by broadcast (multicast or unicast) traffic.

The graph in Figure 12-1 shows broadcast traffic patterns on an interface over a given period of time. The example can also be applied to multicast and unicast traffic. In this example, the broadcast traffic being forwarded exceeded the configured threshold between time intervals T1 and T2 and between T4 and T5. When the amount of specified traffic exceeds the threshold, all traffic of that kind is dropped for the next time period. Therefore, broadcast traffic is blocked during the intervals following T2 and T5. At the next time interval (for example, T3), if broadcast traffic does not exceed the threshold, it is again forwarded.

Figure 12-1 Broadcast Storm Control Example

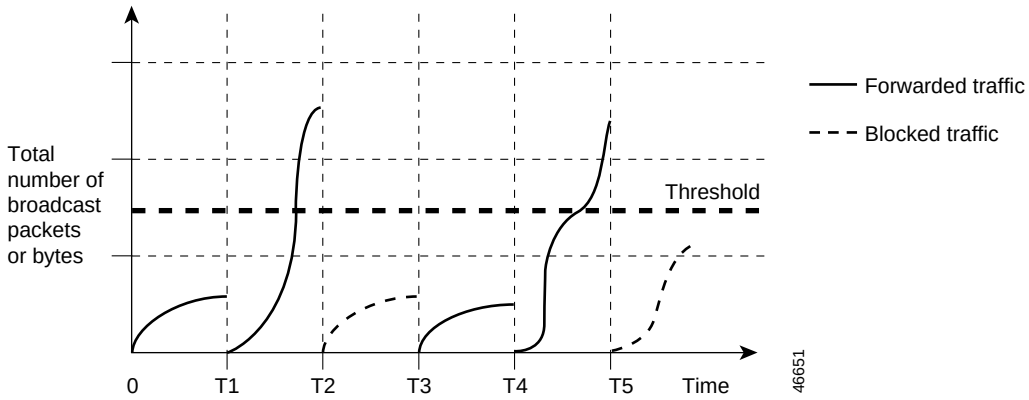

The combination of the storm-control suppression level and the 1-second time interval control the way the storm control algorithm works. A higher threshold allows more packets to pass through. A threshold value of 100 percent means that no limit is placed on the traffic. A value of 0.0 means that all broadcast, multicast, or unicast traffic on that port is blocked.

Note Because packets do not arrive at uniform intervals, the 1-second time interval during which traffic activity is measured can affect the behavior of storm control.

The switch continues to monitor traffic on the port, and when the utilization level is below the threshold level, the type of traffic that was dropped is forwarded again.

You use the **storm-control** interface configuration commands to set the threshold value for each traffic type.

Note Before IOS Release 12.1(8)EA1, you set up storm control threshold values by using the **switchport broadcast**, **switchport multicast**, and **switchport unicast** interface configuration commands. These commands are now obsolete, replaced by the **storm-control** interface configuration commands.

Default Storm Control Configuration

By default, unicast, broadcast, and multicast storm control is disabled on the switch: that is, the suppression level is 100 percent.

Enabling Storm Control

You enable storm control on an interface and enter the percentage of total available bandwidth that you want to be used by a particular type of traffic; entering 100 percent allows all traffic. However, because of hardware limitations and the way in which packets of different sizes are counted, threshold percentages are approximations. Depending on the sizes of the packets making up the incoming traffic, the actual enforced threshold might differ from the configured level by several percentage points.

Note Storm control is supported only on physical interfaces; it is not supported on EtherChannel port channels even though the command is available in the CLI.

Beginning in privileged EXEC mode, follow these steps to enable a particular type of storm control:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the type and number of the physical interface to configure, for example, **gigabitethernet0/1**.
Step 3	**storm-control broadcast level** *level* [.*level*]	Specify the broadcast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all broadcast traffic on that port is blocked.
Step 4	**storm-control multicast level** *level* [.*level*]	Specify the multicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all multicast traffic on that port is blocked.
Step 5	**storm-control unicast level** *level* [.*level*]	Specify the unicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all unicast traffic on that port is blocked.
Step 6	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 7	**show storm-control** [*interface-id*] [**broadcast** | **multicast** | **unicast**]	Verify the storm control suppression levels set on the interface for the specified traffic type. If you do not enter a traffic type, broadcast storm control settings are displayed.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable storm control, use the **no storm-control broadcast level**, **no storm-control multicast level**, or **no storm-control unicast level** interface configuration commands.

This example shows how to set the multicast storm control level at 70.5 percent on Fast Ethernet interface 17 and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface fastethernet0/17
Switch(config-if)# storm-control multicast level 70.5
Switch(config-if)# end
Switch# show storm-control fastethernet0/17 multicast
Interface  Filter State   Level    Current
---------  -------------  -------  -------
Fa0/17     Forwarding      70.50%    0.00%
```

Disabling Storm Control

Beginning in privileged EXEC mode, follow these steps to disable storm control on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the type and number of the physical interface to configure, for example GigabitEthernet 0/1.
Step 3	**no storm-control broadcast level**	Disable broadcast storm control on the interface.
Step 4	**no storm-control multicast level**	Disable multicast storm control on the interface.
Step 5	**no storm-control unicast level**	Disable unicast storm control on the interface.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show storm-control** [*interface-id*] [**broadcast** | **multicast** | **unicast**]	Verify that there are no storm control suppression levels set on the interface for the specified traffic type. If you do not enter a traffic type, broadcast storm control settings are displayed.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

This example shows how to disable the multicast storm control on Fast Ethernet interface 17 and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface fastethernet0/17
Switch(config-if)# no storm-control multicast level
Switch(config-if)# end
Switch# show storm-control fastethernet0/17 multicast
Interface  Filter State   Level    Current
---------  -------------  -------  -------
Fa0/17     inactive       100.00%  N/A
```

Configuring Protected Ports

Some applications require that no traffic be forwarded between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of protected ports ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch.

Protected ports have these features:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.

The default is to have no protected ports defined.

Note The protected port feature is not compatible with fallback bridging. When fallback bridging is enabled, it is possible for packets to be forwarded from one protected port on a switch to another protected port on the same switch if the ports are in different VLANs.

Note There could be times when unknown unicast or multicast traffic from a nonprotected port is flooded to a protected port because a MAC address has timed out or has not been learned by the switch. Use the **switchport block unicast** and **switchport block multicast** interface configuration commands to guarantee that no unicast or multicast traffic is flooded to the port in such a case.

A protected port cannot be a secure port.

You can configure protected ports on a physical interface (for example, Gigabit Ethernet 0/1) or an EtherChannel group (for example, port-channel 5). When you enable protected port for a port channel, it is enabled for all ports in the port channel group.

Beginning in privileged EXEC mode, follow these steps to define a port as a protected port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the type and number of the switchport interface to configure, for example, **gigabitethernet0/1**.
Step 3	**switchport protected**	Configure the interface to be a protected port.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show interfaces** *interface-id* **switchport**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable protected port, use the **no switchport protected** interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/3 as a protected port and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# switchport protected
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/3 switchport
Name: Gi0/3
Switchport: Enabled

<output truncated>

Protected: True
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
```

Configuring Port Blocking

By default, the switch floods packets with unknown destination MAC addresses to all ports. If unknown unicast and multicast traffic is forwarded to a protected port, there could be security issues.

To prevent unknown unicast or multicast traffic from being forwarded from one port to another, you can configure a port (protected or nonprotected) to block unknown unicast or multicast packets.

Note Blocking unicast or multicast traffic is not automatically enabled on protected ports; you must explicitly configure it.

Blocking Flooded Traffic on an Interface

Note The interface can be a physical interface (for example, GigabitEthernet 0/1) or an EtherChannel group (for example, port-channel 5). When you block multicast or unicast traffic for a port channel, it is blocked on all ports in the port channel group.

Beginning in privileged EXEC mode, follow these steps to disable the flooding of multicast and unicast packets to an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the type and number of the switchport interface to configure, for example **gigabitethernet0/1**.
Step 3	**switchport block multicast**	Block unknown multicast forwarding to the port.
Step 4	**switchport block unicast**	Block unknown unicast forwarding to the port.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show interfaces** *interface-id* **switchport**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the interface to the default condition where no traffic is blocked, use the **no switchport block** {**multicast** | **unicast**} interface configuration commands.

This example shows how to block unicast and multicast flooding on Gigabit Ethernet interface 0/1 and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport block multicast
Switch(config-if)# switchport block unicast
Switch(config-if)# end
Switch# show interfaces gigabitethernet0/1 switchport
Name: Gi0/1
Switchport: Enabled

<output truncated>

Protected: True
Unknown unicast blocked: enabled
Unknown multicast blocked: enabled
```

Resuming Normal Forwarding on a Port

Beginning in privileged EXEC mode, follow these steps to resume normal forwarding on a port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode and enter the type and number of the switchport interface to configure, for example **gigabitethernet0/1**.
Step 3	**no switchport block multicast**	Enable unknown multicast flooding to the port.
Step 4	**no switchport block unicast**	Enable unknown unicast flooding to the port.
Step 5	**end**	Return to privileged EXEC mode
Step 6	**show interfaces** *interface-id* **switchport**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Configuring Port Security

You can use the port security feature to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the workstation attached to that port is assured the full bandwidth of the port.

If a port is configured as a secure port and the maximum number of secure MAC addresses is reached, when the MAC address of a station attempting to access the port is different from any of the identified secure MAC addresses, a security violation occurs. Also, if a station with a secure MAC address configured or learned on one secure port attempts to access another secure port, a violation is flagged.

Understanding Port Security

After you have set the maximum number of secure MAC addresses on a port (the range is 1 to 128 with a default of 128), the secure addresses are included in an address table in one of these ways:

- You can configure all secure MAC addresses by using the **switchport port-security mac-address** *mac_address* interface configuration command.
- You can allow the port to dynamically configure secure MAC addresses with the MAC addresses of connected devices.
- You can configure a number of addresses and allow the rest to be dynamically configured.

Note If the port shuts down, all dynamically learned addresses are removed.

Once the maximum number of secure MAC addresses is configured, they are stored in an address table. Setting a maximum number of addresses to one and configuring the MAC address of an attached device ensures that the device has the full bandwidth of the port.

It is a security violation when one of these situations occurs:

- The maximum number of secure MAC addresses have been added to the address table and a station whose MAC address is not in the address table attempts to access the interface.
- A station whose MAC address is configured as a secure MAC address on another secure port attempts to access the interface.

You can configure the interface for one of three violation modes, based on the action to be taken if a violation occurs:

- protect—when the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
- restrict—a port security violation causes a trap notification to be sent to the network management station (NMS).
- shutdown—a port security violation causes a the interface to shut down immediately and an SNMP trap notification is sent. Once shut down, the interface must be manually re-enabled by using the **no shutdown** interface configuration command. This is the default mode.

Default Port Security Configuration

Table 12-1 shows the default port security configuration for an interface.

Table 12-1 Default IGMP Snooping Configuration

Feature	Default Setting
Port security	Disabled on a port
Maximum number of secure MAC addresses	128
Violation mode	Shutdown. The port shuts down when the maximum number of secure MAC addresses is exceeded, and an SNMP trap notification is sent.

Configuration Guidelines

Follow these guidelines when configuring port security:

- A protected port cannot be a routed port.
- A secure port cannot be a dynamic access port or a trunk port.
- A protected port cannot be a secure port.
- A secure port cannot be a destination port for Switch Port Analyzer (SPAN).
- A secure port cannot belong to a Fast EtherChannel or Gigabit EtherChannel port group.
- A secure port cannot be an 802.1X port. If you try to enable 802.1X on a secure port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to a secure port, an error message appears, and the security settings are not changed.

Enabling and Configuring Port Security

Beginning in privileged EXEC mode, follow these steps to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the physical interface to configure, for example **gigabitethernet0/1**.
Step 3	**switchport mode access**	Set the interface mode as access; an interface in the default mode (dynamic desirable) cannot be configured as a secure port.
Step 4	**switchport port-security**	Enable port security on the interface.
Step 5	**switchport port-security maximum** *number of addresses*	(Optional) Set the maximum number of secure MAC addresses for the interface. The range is 1 to 128; the default is 128.

	Command	Purpose
Step 6	**switchport port-security violation** {**protect** | **restrict** | **shutdown**}	(Optional) Set the violation mode, the action to be taken when a security violation is detected, as one of these: • **shutdown**—The interface shuts down immediately, and an SNMP trap notification is sent. When shut down, the interface must be manually re-enabled by using the **no shutdown** interface configuration command. This is the default mode. • **restrict**—A trap notification is sent to the network management station. • **protect**—When the number of port secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
Step 7	**switchport port-security mac-address** *mac_address*	(Optional) Enter a secure MAC address for the interface. You can use this command to enter the maximum number of secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned.
Step 8	**end**	Return to privileged EXEC mode.
Step 9	**show port-security interface** *interface-id* **show port-security address**	Verify your entries.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the interface to the default condition as not a secure port, use the **no switchport port-security** interface configuration command.

To return the interface to the default number of secure MAC addresses (128), use the no **switchport port-security maximum** *number of addresses.*

To delete a MAC address from the address table, use the **no switchport port-security mac-address** *mac_address* command.

To return the violation mode to the default condition (shutdown mode), use the **no switchport port-security violation** {**protocol** | **restrict**} command.

This example shows how to enable port security on Fast Ethernet port 12 and to set the maximum number of secure addresses to 5. The violation mode is the default, and no secure MAC addresses are configured.

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface fastethernet0/12
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 5
Switch(config-if)# end
Switch# show port-security interface fastethernet0/12
Security Enabled:Yes, Port Status:SecureUp
Violation Mode:Shutdown
Max. Addrs:5, Current Addrs:0, Configure Addrs:0
```

This example shows how to configure a secure MAC address on Fast Ethernet port 12 and verify the configuration.

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface fastethernet0/12
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security mac-address 1000.2000.3000
Switch(config-if)# end
Switch# show port-security address
          Secure Mac Address Table
------------------------------------------------------------

Vlan    Mac Address       Type                Ports
----    -----------       ----                -----
   1    1000.2000.3000    SecureConfigured     Fa0/12
```

Displaying Port-Based Traffic Control Settings

The **show interfaces** *interface-id* **switchport** privileged EXEC command displays (among other characteristics) the interface traffic suppression and control configuration. The **show interfaces counters** privileged EXEC commands display the count of discarded packets. The **show storm control** and **show port-security** privileged EXEC commands display those features.

To display traffic control information, use one or more of the privileged EXEC commands in Table 12-2.

Table 12-2 Commands for Displaying Traffic Control Status and Configuration

Command	Purpose
show interfaces [*interface-id*] **switchport**	Displays the administrative and operational status of all switching (nonrouting) ports or the specified port, including port blocking and port protection settings.
show storm-control [*interface-id*] [**broadcast** | **multicast** | **unicast**]	Displays storm control suppression levels set on all interfaces or the specified interface for the specified traffic type or for broadcast traffic if no traffic type is entered.
show interfaces [*interface-id*] **counters broadcast**	Displays the storm-control broadcast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [*interface-id*] **counters multicast**	Displays the storm-control multicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [*interface-id*] **counters unicast**	Displays the storm-control unicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show port-security [**interface** *interface-id*]	Displays port security settings for the switch or for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode.
show port-security [**interface** *interface-id*] **address**	Displays all secure MAC addresses configured on all switch interfaces or on a specified interface.

This is a an example of output from the **show interfaces switchport** privileged EXEC command:

```
Switch# show interfaces gigabitethernet0/1 switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: False
Unknown unicast blocked: disabled
Unknown multicast blocked: enabled
```

This is a an example of output from the **show interfaces counters broadcast** privileged EXEC command:

```
Switch# show interfaces counters broadcast

Port      BcastSuppDiscards
Gi0/1                     0
Gi0/2                     0
Gi0/3                     0
Gi0/4                     0
```

This is an example of output from the **show switchport port-security** privileged EXEC command when you do not enter an interface.

```
Switch# show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                    (Count)        (Count)      (Count)
-------------------------------------------------------------------------------
     Fa0/12          25            1             0            Restrict
     Fa0/22          15            1             0            Protect
```

This is an example of output from the **show switchport port-security** privileged EXEC command for a specified interface.

```
Switch# show port-security interface fastethernet0/12
Security Enabled:Yes, Port Status:SecureUp
Violation Mode:Restrict
Max. Addrs:25, Current Addrs:1, Configure Addrs:1
```

This is a an example of output from the **show port-security address** privileged EXEC command.

```
Switch# show port-security address
          Secure Mac Address Table
------------------------------------------------------------

Vlan    Mac Address       Type                Ports
----    -----------       ----                -----
   1    1000.2000.3000    SecureConfigured     Fa0/12
```

This is an example of output from the **show storm-control** command when no keywords are entered. Because no traffic type keyword was entered, the broadcast storm control settings are displayed.

```
Switch# show storm-control
Interface  Filter State   Level    Current
---------  -------------  -------  -------
Fa0/1      inactive       100.00%  N/A
Fa0/2      inactive       100.00%  N/A
Fa0/3      inactive       100.00%  N/A
Fa0/4      inactive       100.00%  N/A
Fa0/5      inactive       100.00%  N/A
Fa0/6      inactive       100.00%  N/A
Fa0/7     Forwarding      50.00%    0.00%
Fa0/8      inactive       100.00%  N/A

<output truncated>
```

This is an example of output from the **show storm-control** command for a specified interface. Because no traffic type keyword was entered, the broadcast storm control settings are displayed.

```
Switch# show storm-control fastethernet0/17
Interface  Filter State   Level    Current
---------  -------------  -------  -------
Fa0/17     Forwarding      50.00%    0.00%
```

This is an example of output from the **show storm-control** command for a specified interface and traffic type, where no storm control threshold has been set for that traffic type on the specified interface.

```
Switch# show storm-control fastethernet0/17 multicast
Interface  Filter State   Level    Current
---------  -------------  -------  -------
Fa0/17     inactive       100.00%  N/A
```

CHAPTER 13

Configuring CDP

This chapter describes how to configure Cisco Discovery Protocol (CDP) on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding CDP, page 13-1
- Configuring CDP, page 13-2
- Monitoring and Maintaining CDP, page 13-5

Understanding CDP

CDP is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

On the Catalyst 3550 switch, CDP enables the Cluster Management Suite to display a graphical view of the network. The switch uses CDP to find cluster candidates and maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch by default.

The switch supports CDP version 2.

Configuring CDP

These sections include CDP configuration information and procedures:

- Default CDP Configuration, page 13-2
- Configuring the CDP Characteristics, page 13-2
- Disabling and Enabling CDP, page 13-3
- Disabling and Enabling CDP on an Interface, page 13-4

Default CDP Configuration

Table 13-1 shows the default CDP configuration.

Table 13-1 Default CDP Configuration

Feature	Default Setting
CDP global state	Enabled
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP version-2 advertisements	Enabled

Configuring the CDP Characteristics

You can configure the frequency of CDP updates, the amount of time to hold the information before discarding it, and whether or not to send version-2 advertisements.

Beginning in privileged EXEC mode, follow these steps to configure the CDP timer, holdtime, and advertisement type.

Note Steps 2 through 4 are all optional and can be performed in any order.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**cdp timer** *seconds*	(Optional) Set the transmission frequency of CDP updates in seconds. The range is from 5 to 254; the default is 60 seconds.
Step 3	**cdp holdtime** *seconds*	(Optional) Specify the amount of time a receiving device should hold the information sent by your device before discarding it. The range is from 10 to 255 seconds; the default is 180 seconds.
Step 4	**cdp advertise-v2**	(Optional) Configure CDP to send version-2 advertisements. This is the default state.
Step 5	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 6	**show cdp**	Verify configuration by displaying global information about CDP on the device.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** form of the CDP commands to return to the default settings.

This example shows how to configure and verify CDP characteristics.

```
Switch# configure terminal
Switch(config)# cdp timer 50
Switch(config)# cdp holdtime 120
Switch(config)# cdp advertise-v2
Switch(config)# end

Switch# show cdp
Global CDP information:
        Sending CDP packets every 50 seconds
        Sending a holdtime value of 120 seconds
        Sending CDPv2 advertisements is enabled
```

For additional CDP **show** commands, see the "Monitoring and Maintaining CDP" section on page 13-5.

Disabling and Enabling CDP

CDP is enabled by default.

Note Creating and maintaining switch clusters is based on the regular exchange of CDP messages. Disabling CDP can interrupt cluster discovery. For more information, see Chapter 5, "Clustering Switches."

Beginning in privileged EXEC mode, follow these steps to disable the CDP device discovery capability:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no cdp run**	Disable CDP.
Step 3	**end**	Return to privileged EXEC mode.

Beginning in privileged EXEC mode, follow these steps to enable CDP when it has been disabled:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**cdp run**	Enable CDP after disabling it.
Step 3	**end**	Return to privileged EXEC mode.

This example shows how to enable CDP if it has been disabled.

```
Switch# configure terminal
Switch(config)# cdp run
Switch(config)# end
```

Disabling and Enabling CDP on an Interface

CDP is enabled by default on all supported interfaces to send and receive CDP information.

Beginning in privileged EXEC mode, follow these steps to disable CDP on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the interface on which you are disabling CDP.
Step 3	**no cdp enable**	Disable CDP on an interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Beginning in privileged EXEC mode, follow these steps to enable CDP on an interface when it has been disabled:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the interface on which you are enabling CDP.
Step 3	**cdp enable**	Enable CDP on an interface after disabling it.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

This example shows how to enable CDP on an interface when it has been disabled.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/5
Switch(config-if)# cdp enable
Switch(config-if)# end
```

Monitoring and Maintaining CDP

To monitor and maintain CDP on your device, perform one or more of these tasks, beginning in privileged EXEC mode.

Command	Description
clear cdp counters	Reset the traffic counters to zero.
clear cdp table	Delete the CDP table of information about neighbors.
show cdp	Display global information, such as frequency of transmissions and the holdtime for packets being sent.
show cdp entry *entry-name* [**protocol** | **version**]	Display information about a specific neighbor. You can enter an asterisk (*) to display all CDP neighbors, or you can enter the name of the neighbor about which you want information. You can also limit the display to information about the protocols enabled on the specified neighbor or information about the version of software running on the device.
show cdp interface [*type number*]	Display information about interfaces where CDP is enabled. You can limit the display to the type of interface or the number of the interface about which you want information (for example, entering **gigabitethernet 0/1** displays information only about Gigabit Ethernet port 1).
show cdp neighbors [*type number*] [**detail**]	Display information about neighbors, including device type, interface type and number, holdtime settings, capabilities, platform, and port ID. You can limit the display to neighbors on a specific type or number of interface or expand the display to provide more detailed information.
show cdp traffic	Display CDP counters, including the number of packets sent and received and checksum errors.

These are examples of outputs from the CDP **show** privileged EXEC commands:

```
Switch# show cdp
Global CDP information:
        Sending CDP packets every 50 seconds
        Sending a holdtime value of 120 seconds
        Sending CDPv2 advertisements is enabled
```

```
Switch# show cdp entry *
-------------------------
Device ID: Switch
Entry address(es):
  IP address: 10.1.1.66
Platform: cisco WS-C3550-12T,  Capabilities: Switch IGMP
Interface: GigabitEthernet0/2,  Port ID (outgoing port): GigabitEthernet0/2
Holdtime : 129 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) C3550 Software (C3550-I5Q3L2-M), Experimental Version 12.1(20010612:021
316) [jang-flamingo 120]
Copyright (c) 1986-2001 by cisco Systems, Inc.
Compiled Fri 06-Jul-01 18:18 by jang

advertisement version: 2
Protocol Hello:  OUI=0x00000C, Protocol ID=0x0112; payload len=27, value=0000000
0FFFFFFFF010221FF0000000000000000024B293A00FF0000
VTP Management Domain: ''
Duplex: full

-------------------------
Device ID: idf2-1-lab-l3.cisco.com
Entry address(es):
  IP address: 10.1.1.10
Platform: cisco WS-C3524-XL,  Capabilities: Trans-Bridge Switch
Interface: GigabitEthernet0/1,  Port ID (outgoing port): FastEthernet0/10
Holdtime : 141 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) C3500XL Software (C3500XL-C3H2S-M), Version 12.0(5.1)XP, MAINTENANCE IN
TERIM SOFTWARE
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Fri 10-Dec-99 11:16 by cchang

advertisement version: 2
Protocol Hello:  OUI=0x00000C, Protocol ID=0x0112; payload len=25, value=0000000
0FFFFFFFF010101FF0000000000000000142EFA400FF
VTP Management Domain: ''

Switch# show cdp entry * protocol
Protocol information for talSwitch14 :
  IP address: 172.20.135.194
Protocol information for tstswitch2 :
  IP address: 172.20.135.204
  IP address: 172.20.135.202
Protocol information for tstswitch2 :
  IP address: 172.20.135.204
  IP address: 172.20.135.202
```

```
Switch# show cdp interface
GigabitEthernet0/1 is up, line protocol is up
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/2 is up, line protocol is down
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/3 is administratively down, line protocol is down
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/4 is up, line protocol is down
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/5 is up, line protocol is up
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/6 is up, line protocol is up
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/7 is up, line protocol is down
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds
GigabitEthernet0/8 is up, line protocol is down
  Encapsulation ARPA
  Sending CDP packets every 60 seconds
  Holdtime is 180 seconds

Switch# show cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce       Holdtme      Capability     Platform       Port ID
Perdido2         Gig 0/6             125          R S I          WS-C3550-1Gig      0/6
Perdido2         Gig 0/5             125          R S I          WS-C3550-1Gig      0/5

Switch# show cdp traffic
CDP counters :
        Total packets output: 50882, Input: 52510
        Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
        No memory: 0, Invalid packet: 0, Fragmented: 0
        CDP version 1 advertisements output: 0, Input: 0
        CDP version 2 advertisements output: 50882, Input: 52510
```

CHAPTER 14

Configuring UDLD

This chapter describes how to configure the UniDirectional Link Detection (UDLD) protocol on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding UDLD, page 14-1
- Configuring UDLD, page 14-3
- Displaying UDLD Status, page 14-5

Understanding UDLD

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it administratively shuts down the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by the local device is received by the neighbor but traffic from the neighbor is not received by the local device. If one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally from a Layer 1 perspective, UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation because autonegotiation operates at Layer 1.

UDLD operates by using two mechanisms:

- Neighbor database maintenance

  UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors.

  When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

  Whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset, UDLD clears all existing cache entries for the interfaces affected by the configuration change. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

  UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply. If the detection window ends and no valid reply message is received, the link is considered unidirectional, and the interface is shut down.

Figure 14-1 shows an example of a unidirectional link condition.

Figure 14-1 UDLD Detection of a Unidirectional Link

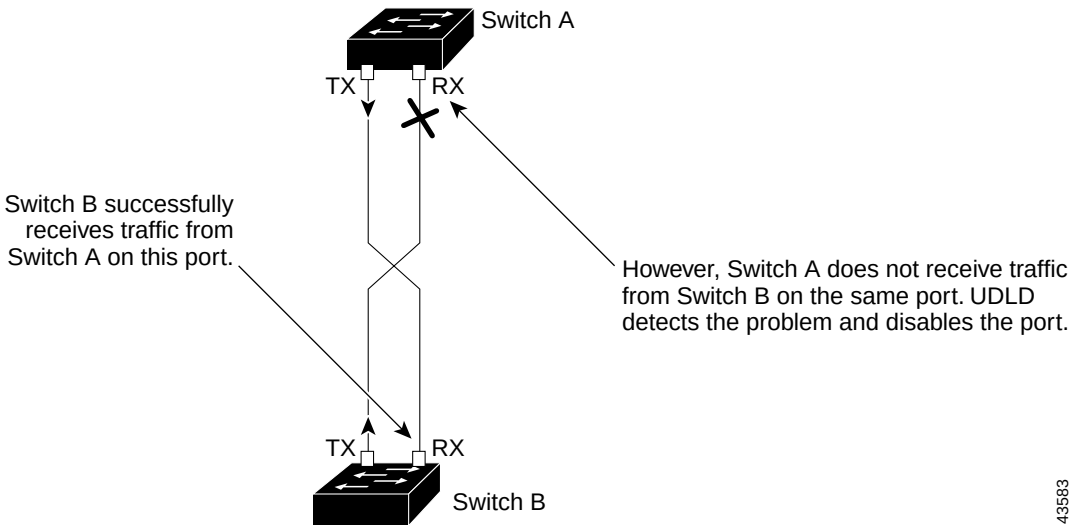
Switch A
TX
RX
Switch B successfully receives traffic from Switch A on this port.
However, Switch A does not receive traffic from Switch B on the same port. UDLD detects the problem and disables the port.
TX
RX
Switch B
43583

Configuring UDLD

This section describes how to configure UDLD on your switch. It contains this configuration information:

- Default UDLD Configuration, page 14-3
- Enabling UDLD Globally, page 14-3
- Enabling UDLD on an Interface, page 14-4
- Resetting an Interface Shut Down by UDLD, page 14-4

Default UDLD Configuration

Table 14-1 shows the default UDLD configuration.

Table 14-1 Default UDLD Configuration

Feature	Default Setting
UDLD global enable state	Globally disabled
UDLD per-interface enable state for fiber-optic media	Disabled on all Ethernet fiber-optic interfaces
UDLD per-interface enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX interfaces

UDLD is not supported on ATM interfaces. A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.

Enabling UDLD Globally

Beginning in privileged EXEC mode, follow these steps to enable UDLD globally on all fiber-optic interfaces on the switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**udld enable**	Enable UDLD on all fiber-optic interfaces on the switch. UDLD is disabled by default. This command affects fiber-optic interfaces only. Use the **udld** interface configuration command to enable UDLD on other interface types. For more information, see the "Enabling UDLD on an Interface" section on page 14-4. An individual interface configuration overrides the setting of the **udld enable** global configuration command.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show udld**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable UDLD globally on fiber-optic interfaces, use the **no udld enable** global configuration command.

Enabling UDLD on an Interface

Beginning in privileged EXEC mode, follow these steps to enable UDLD on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be enabled for UDLD.
Step 3	**udld enable**	Enable UDLD on the specified interface. On a fiber-optic interface, this command overrides the **udld enable** global configuration command setting.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show udld** *interface-id*	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable UDLD on a non-fiber-optic interface, use the **no udld enable** interface configuration command.

Note On fiber-optic interfaces, the **no udld enable** command reverts the interface configuration to the **udld enable** global configuration command setting.

To disable UDLD on a fiber-optic interface, use the **udld disable** command to revert to the **udld enable** global configuration command setting. This command is not supported on non-fiber-optic interfaces.

Resetting an Interface Shut Down by UDLD

Beginning in privileged EXEC mode, follow these steps to reset all interfaces shut down by UDLD:

	Command	Purpose
Step 1	**udld reset**	Reset all interfaces shut down by UDLD.
Step 2	**show udld**	Verify your entries.
Step 3	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

You can also bring up the interface by using these commands:

- The **no shutdown** interface configuration command restarts the disabled interface.
- The **no udld enable** global configuration command re-enables UDLD globally.
- The **udld disable** interface configuration command re-enables UDLD on the specified interface.

Displaying UDLD Status

To display the UDLD status for the specified interface or for all interfaces, use the **show udld** [*interface-id*] privileged EXEC command.

Configuring SPAN

This chapter describes how to configure Switch Port Analyzer (SPAN) on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding SPAN, page 15-1
- Configuring SPAN, page 15-6
- Displaying SPAN Status, page 15-13

Understanding SPAN

You can analyze network traffic passing through ports or VLANs by using SPAN to send a copy of the traffic to another port on the switch that has been connected to a SwitchProbe device or other Remote Monitoring (RMON) probe. SPAN mirrors received or sent (or both) traffic on a source port, or received traffic on one or more source ports or source VLANs, to a destination port for analysis.

For example, in Figure 15-1, all traffic on Gigabit Ethernet port 5 (the source port) is mirrored to Gigabit Ethernet port 10 (the destination port). A network analyzer on port 10 receives all network traffic from port 5 without being physically attached to port 5.

Figure 15-1 Example SPAN Configuration

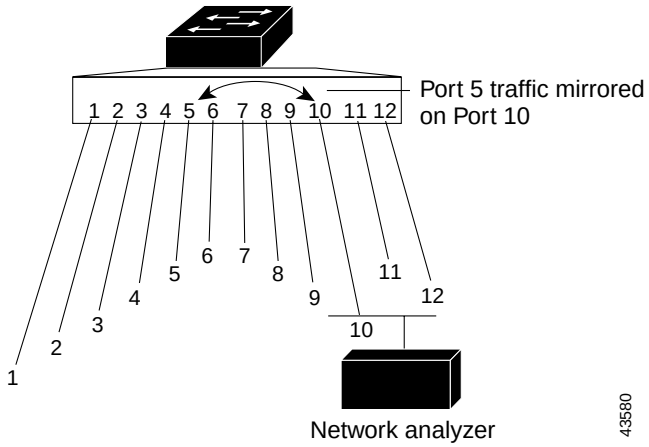

Only traffic that enters or leaves source ports or traffic that enters source VLANs can be monitored by using SPAN; traffic that gets routed to ingress source ports or source VLANs cannot be monitored. For example, if incoming traffic is being monitored, traffic that gets routed from another VLAN to the source VLAN is not monitored; however, traffic that is received on the source VLAN and routed to another VLAN is monitored.

This release supports only local SPAN, which means the source and destination interfaces must be on the same switch.

SPAN does not affect the switching of network traffic on source ports or source VLANs; a copy of the packets received or sent by the source interfaces are sent to the destination interface. However, an oversubscribed SPAN destination, for example, a 10-Mbps port monitoring a 100-Mbps port, can cause congestion on the switch. Destination ports do not receive or forward traffic, except that required for the SPAN session.

SPAN Concepts and Terminology

This section describes concepts and terminology associated with SPAN configuration.

SPAN Session

A SPAN session is an association of a destination port with source ports or source VLANs. You configure SPAN sessions by using parameters that specify the type of network traffic to monitor. Traffic monitoring in a SPAN session has these restrictions:

- You can monitor incoming traffic on a series or range of ports or VLANs.
- You can monitor outgoing traffic on a single port; you cannot monitor outgoing traffic on multiple ports.
- You cannot monitor outgoing traffic on VLANs.

You can configure two separate SPAN sessions with separate or overlapping sets of SPAN source VLANs. Both switched and routed ports can be configured as SPAN sources and destinations. SPAN sessions do not interfere with the normal operation of the switch.

You can configure SPAN sessions on disabled ports; however, a SPAN session does not become active unless you enable the destination port and at least one source port or VLAN for that session. The **show monitor session** *session_number* privileged EXEC command displays the operational status of a SPAN session.

A SPAN session remains inactive after system power-on until the destination port is operational.

Traffic Types

SPAN sessions include these traffic types:

- Receive (Rx) SPAN—The goal of receive (or ingress) SPAN is to monitor as much as possible all the packets received by the source interface or VLAN before any modification or processing is performed by the switch. A copy of each packet received by the source is sent to the destination port for that SPAN session. You can monitor a series or range of ingress ports or VLANs in a SPAN session.

  On tagged packets (ISL or 802.1Q), the tagging is removed at the ingress port. At the destination port, if tagging is enabled, the packets are seen with the ISL or 802.1Q headers, as specified. If no tagging is specified, packets are seen in the native format.

  Packets that are modified because of routing are copied without modification for Rx SPAN; that is, the original packet is copied. Packets that are modified because of quality of service (QoS)—for example, modified Differentiated Services Code Point (DSCP)—are copied with modification for Rx SPAN.

  Some features that can cause a packet to be dropped during receive processing have no effect on SPAN; the destination port receives a copy of the packet even if the actual incoming packet is dropped. These features include IP standard and extended input access control lists (ACLs), IP standard and extended output ACLs for unicast, VLAN maps, ingress QoS policing, and policy-based routing. Switch congestion that causes packets to be dropped also has no effect on SPAN.

- Transmit (Tx) SPAN—The goal of transmit (or egress) SPAN is to monitor as much as possible all the packets sent by the source interface after all modification and processing is performed by the switch. A copy of each packet sent by the source is sent to the destination port for that SPAN session. The copy is provided after the packet is modified.

  Only one egress source port is allowed in one SPAN session. VLAN monitoring is not supported in the egress direction.

  Packets that are modified because of routing—for example, with a time-to-live (TTL) or MAC-address modification—are duplicated at the destination port. On packets that are modified because of QoS, the modified packet might not have the same DSCP (IP packet) or CoS (non-IP packet) as the SPAN source.

  Some features that can cause a packet to be dropped during transmit processing might also affect the duplicated copy for SPAN. These features include VLAN maps, IP standard and extended output ACLs on multicast packets, and egress QoS policing. In the case of output ACLs, if the SPAN source drops the packet, the SPAN destination would also drop the packet. In the case of egress QoS policing, if the SPAN source drops the packet, the SPAN destination might not drop it. If the source port is oversubscribed, the destination ports will have different dropping behavior.

- Both—In a SPAN session, a single port can be monitored for both received and sent packets.

Source Port

A source port (also called a *monitored port*) is a switched or routed port that you monitor for network traffic analysis. In a single SPAN session, you can monitor source port traffic such as received (Rx), transmitted (Tx), or bidirectional (both); however, on a VLAN, you can monitor only received traffic. The switch supports any number of source ports (up to the maximum number of available ports on the switch) and any number of source ingress VLANs (up to the maximum number of VLANs supported).

A source port has these characteristics:

- It can be any port type (for example, EtherChannel, Fast Ethernet, Gigabit Ethernet, and so forth).
- It can be monitored in multiple SPAN sessions.
- It cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For EtherChannel sources, the monitored direction would apply to all the physical ports in the group.
- Source ports can be in the same or different VLANs.

For VLAN SPAN (VSPAN), all active ports in the source VLAN are included as source ports.

You can configure a trunk port as a source port. By default, all VLANs active on the trunk are monitored on a trunk source port. You can limit SPAN traffic monitoring on trunk source ports to specific VLANs by using trunk VLAN filtering, which is the analysis of network traffic on a selected set of VLANs on source trunk ports. Only switched traffic in the selected VLANs is sent to the destination port. This feature affects only traffic forwarded to the destination SPAN port and does not affect the switching of normal traffic. This feature is not applicable for VLAN SPAN sessions.

Destination Port

Each SPAN session must have a destination port (also called a *monitoring port*) that receives a copy of traffic from the source port.

The destination port has these characteristics:

- It must reside on the same switch as the source port.
- It can be any Ethernet physical port.
- It can participate in only one SPAN session at a time (a destination port in one SPAN session cannot be a destination port for a second SPAN session).
- It cannot be a source port.
- It cannot be an EtherChannel port or a VLAN.
- When it is active, incoming traffic is disabled; it does not forward any traffic except that required for the SPAN session.
- It does not participate in spanning tree while the SPAN session is active.
- When it is an active destination port, it does not participate in any of the Layer 2 protocols (STP, VTP, CDP, DTP, PagP).
- A destination port that belongs to a source VLAN of any SPAN session is excluded from the source list and is not monitored.
- No address learning occurs on the destination port.

VLAN-Based SPAN

VLAN-based SPAN (VSPAN) is the analysis of the network traffic in one or more VLANs. You can configure VSPAN to monitor only received (Rx) traffic, which applies to all the ports for that VLAN.

Use these guidelines for VSPAN sessions:

- Trunk ports are included as source ports for VSPAN sessions.
- Only traffic with the monitored VLAN number is sent to the destination port.
- If a destination port belongs to a source VLAN, it is excluded from the source list and is not monitored.
- If ports are added to or removed from the source VLANs, they are added to or removed from the source ports being monitored.
- VLAN pruning and the VLAN allowed list have no effect on SPAN monitoring.
- VSPAN only monitors traffic that enters the switch, not traffic that is routed between VLANs. For example, if a VLAN is being Rx-monitored and the multilayer switch routes traffic from another VLAN to the monitored VLAN, that traffic is not monitored and is not received on the SPAN destination port.

SPAN Traffic

You can use SPAN to monitor all network traffic, including multicast and bridge protocol data unit (BPDU) packets, and CDP, VTP, DTP, STP, and PagP packets. Multicast packet monitoring is enabled by default.

In some SPAN configurations, multiple copies of the same source packet are sent to the SPAN destination port. For example, a bidirectional (both Rx and Tx) SPAN session is configured for sources a1 and a2 to a destination port d1. If a packet enters the switch through a1 and is switched to a2, both incoming and outgoing packets are sent to destination port d1. Both packets are the same (unless a Layer-3 rewrite occurs, in which case the packets are different because of the added Layer 3 information).

SPAN Interaction with Other Features

SPAN interacts with these features:

- Routing—Ingress SPAN does not monitor routed traffic. VSPAN only monitors traffic that enters the switch, not traffic that is routed between VLANs. For example, if a VLAN is being Rx-monitored and the multilayer switch routes traffic from another VLAN to the monitored VLAN, that traffic is not monitored and not received on the SPAN destination port.
- Spanning Tree Protocol (STP)—A destination port does not participate in STP while its SPAN session is active. The destination port can participate in STP after the SPAN session is disabled. On a source port, SPAN does not affect the STP status.

Caution Make sure there are no potential loops in the network topology when you enable incoming traffic for a destination port.

- Cisco Discovery Protocol (CDP)—A SPAN destination port does not participate in CDP while the SPAN session is active. After the SPAN session is disabled, the port again participates in CDP

- VLAN and trunking—You can modify VLAN membership or trunk settings for source and destination ports at any time. However, changes in VLAN membership or trunk settings for a destination port do not take effect until you disable the SPAN session. Changes in VLAN membership or trunk settings for a source port immediately take effect, and the respective SPAN sessions automatically adjust accordingly.
- EtherChannel—You can configure an EtherChannel group as a source port but not as a SPAN destination port. SPAN configuration fails if the destination port is part of an EtherChannel group. When a channel group is configured as a SPAN source, the entire group is monitored.

  If a port is added to a monitored EtherChannel group, the new port is added to the SPAN source port list. If a port is removed from a monitored EtherChannel group, it is automatically removed from the source port list. If the port is the only port in the EtherChannel group, the EtherChannel group is removed from SPAN.

  If a physical port that belongs to an EtherChannel group is configured as a SPAN source or destination port, it is removed from the EtherChannel group. After the port is removed from the SPAN session, it rejoins the EtherChannel group.
- QoS—For ingress monitoring, the packets sent to the SPAN destination port might be different from the packets actually received at the SPAN source port because the packets are forwarded after ingress QoS classification and policing. The packet DSCP might not be the same as the received packet.

  For egress monitoring, the packets sent out the SPAN destination port might not be the same as the packets sent out of SPAN source ports because the egress QoS policing at the SPAN source port might change the packet classification. QoS policing is not applied at SPAN destination ports.
- Multicast traffic can be monitored. For egress and ingress port monitoring, only a single unedited packet is sent to the SPAN destination port. It does not reflect the number of times the multicast packet is sent.
- A secure port cannot be a SPAN destination port.
- You can enable 802.1X on a port that is a SPAN destination port; however, 802.1X is disabled until the port is removed as a SPAN destination. You can enable 802.1X on a SPAN source port.

Configuring SPAN

This section describes how to configure SPAN on your switch. It contains this configuration information:

- Default SPAN Configuration, page 15-7
- SPAN Configuration Guidelines, page 15-7
- Creating a SPAN Session and Specifying Ports to Monitor, page 15-8
- Removing Ports from a SPAN Session, page 15-10
- Specifying VLANs to Monitor, page 15-11
- Specifying VLANs to Filter, page 15-12

Default SPAN Configuration

Table 15-1 shows the default SPAN configuration.

This release supports only local SPAN; remote SPAN (RSPAN) is not supported.

Table 15-1 Default SPAN Configuration

Feature	Default Setting
SPAN state	Disabled
Source port traffic to monitor	Both received and sent traffic (both) **Note** Only received traffic can be monitored on source VLANs or multiple source ports
Encapsulation type (destination port)	Native form (no encapsulation type header)

SPAN Configuration Guidelines

Follow these guidelines when configuring SPAN:

- Use a network analyzer to monitor ports.
- Only two SPAN sessions can be active on a switch at the same time.
- The destination port cannot be a source port; a source port cannot be a destination port.
- You can have only one destination port per SPAN session. You cannot have two SPAN sessions using the same destination port.
- An EtherChannel port can be a SPAN source port; it cannot be a SPAN destination port.
- For a SPAN source port, you can monitor transmitted traffic for a single port or received traffic for a series or range of ports or VLANs.
- When you configure a switch port as a SPAN destination port, it is no longer a normal switch port; only monitored traffic passes through the SPAN destination port.
- A trunk port can be a source port or a destination port. When a destination port is a trunk port, outgoing packets through the SPAN port carry the encapsulation headers configured by the user—either Inter-Switch Link (ISL) or IEEE 802.1Q. If no encapsulation type is defined, the packets are sent in native form.
- When you specify a single source port and do not specify a traffic type (Tx, Rx, or both), both is used as the default.
- You can configure a disabled port to be a source or destination port, but the SPAN function does not start until the destination port and at least one source port or source VLAN are enabled.
- For received traffic, you can mix multiple source port and source VLANs within a single SPAN session. You cannot mix source VLANs and filter VLANs within a SPAN session; you can have source VLANs or filter VLANs, but not both at the same time.
- You can limit SPAN traffic to specific VLANs by using the **filter vlan** keyword. If a trunk port is being monitored, only traffic on the VLANs specified with this keyword is monitored. By default, all VLANs are monitored on a trunk port.

- The **no monitor session** *session_number* global configuration command removes a source or destination port from the SPAN session or removes a source VLAN from the SPAN session. If you do not specify any options following the **no monitor session** *session_number* command, the entire SPAN session is removed. The **no monitor** global configuration command also clears all SPAN sessions.
- A SPAN destination port never participates in any VLAN spanning tree. SPAN does include BPDUs in the monitored traffic, so any spanning-tree BPDUs received on the SPAN destination port for a SPAN session were copied from the SPAN source ports.
- When SPAN is enabled, configuration changes have these results:
  - If you change the VLAN configuration of a destination port, the change is not effective until SPAN is disabled.
  - If you disable all source ports or the destination port, the SPAN function stops until both a source and destination port are enabled.
  - If the source is a VLAN, the number of ports being monitored changes when you move a switched port in or out of the monitored VLAN.

Creating a SPAN Session and Specifying Ports to Monitor

Beginning in privileged EXEC mode, follow these steps to create a SPAN session and specify the source (monitored) and destination (monitoring) ports:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no monitor session** *session_number*	Clear any existing SPAN configuration for the session.
Step 3	**monitor session** *session_number* **source interface** *interface-id* [**,** | **-**] [**both** | **rx** | **tx**]	Specify the SPAN session and the source port (monitored port). For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) [**,** | **-**] —Specify a series or range of interfaces. This option is valid when monitoring only received traffic. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both transmitted and received traffic. • **both**—Monitor both received and transmitted traffic. • **rx**—Monitor received traffic. • **tx**—Monitor transmitted traffic.

	Command	Purpose
Step 4	**monitor session** *session_number* **destination interface** *interface-id* [**encapsulation** {**dot1q** | **isl**}]	Specify the SPAN session and the destination port (monitoring port). For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. • **isl**—Use ISL encapsulation. • **dot1q**—Use 802.1Q encapsulation.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show monitor** [**session** *session_number*]	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the entire SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a source or destination port from the SPAN session, use the **no monitor session** *session_number* **source interface** *interface-id* global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command.

This example shows how to set up a SPAN session, session 1, for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is cleared, and then bidirectional traffic is mirrored from source port 1 to destination port 10. The **show monitor session** privileged EXEC command is used to verify the configuration.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/1
Switch(config)# monitor session 1 destination interface gigabitethernet0/10 encapsulation
dot1q
Switch(config)# end
Switch# show monitor session 1
Session 1
---------
Source Ports:
    RX Only:       None
    TX Only:       None
    Both:          Gi0/1
Source VLANs:
    RX Only:       None
    TX Only:       None
    Both:          None
Destination Ports: Gi0/10
    Encapsulation: DOT1Q
Filter VLANs:      None
```

Removing Ports from a SPAN Session

Beginning in privileged EXEC mode, follow these steps to remove a port as a SPAN source for a session:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no monitor session** *session_number* **source interface** *interface-id* [**,** | **-**] [**both** | **rx** | **tx**]	Specify the characteristics of the source port (monitored port) and SPAN session to remove. For *session*, specify 1 or 2. For *interface-id*, specify the source port to no longer monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (**port-channel** *port-channel-number*). (Optional) Use [**,** | **-**] to specify a series or range of interfaces if they were configured. This option is valid when monitoring only received traffic. Enter a space after the comma; enter a space before and after the hyphen. (Optional) Specify the direction of traffic (**both**, **rx**, or **tx**) to no longer monitor. If you do not specify a traffic direction, both transmit and receive are disabled.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show monitor** [**session** *session_number*]	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a destination port from the SPAN session, use the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. To change the encapsulation type back to the default (native), use the **monitor session** *session_number* **destination interface** *interface-id* without the **encapsulation** keyword.

This example shows how to remove port 1 as a SPAN source for SPAN session 1 and verify the configuration:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1
Switch(config)# end
Switch# show monitor session 1
Session 1
---------
Source Ports:
    RX Only:       None
    TX Only:       None
    Both:          None
Source VLANs:
    RX Only:       None
    TX Only:       None
    Both:          None
Destination Ports:Gi0/6
    Encapsulation:DOT1Q
Filter VLANs:     None
```

This example shows how to disable received traffic monitoring on port 1, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1 rx
```

The monitoring of traffic received on port 1 is disabled, but traffic sent from this port continues to be monitored.

Specifying VLANs to Monitor

VLAN monitoring is similar to port monitoring. Beginning in privileged EXEC mode, follow these steps to specify VLANs to monitor:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no monitor session** *session_number*	Clear any existing SPAN configuration for the session. For *session_number*, specify 1 or 2.
Step 3	**monitor session** *session_number* **source vlan** *vlan-id* [**,** | **-**] **rx**	Specify the SPAN session and the source VLANs (monitored VLANs). You can monitor only received (**rx**) traffic on VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 1005. (Optional) Use a comma (**,**) to specify a series of VLANs, or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
Step 4	**monitor session** *session_number* **destination interface** *interface-id* [**encapsulation** {**dot1q** | **isl**}]	Specify the SPAN session and the destination port (monitoring port). For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces. (Optional) Specify the encapsulation header for outgoing packets. If not specified, packets are sent in native form. • **isl**—Use ISL encapsulation. • **dot1q**—Use 802.1Q encapsulation.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show monitor** [**session** *session_number*]	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove one or more source VLANs or destination ports from the SPAN session, use the **no monitor session** *session_number* **source vlan** *vlan-id* global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command.

This example shows how to clear any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on all ports belonging to VLANs 1 through 3, and send it to destination port 7. The configuration is then modified to also monitor received traffic on all ports belonging to VLAN 10.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source vlan 1 - 3 rx
Switch(config)# monitor session 2 destination interface gigabitethernet0/7
Switch(config)# monitor session 2 source vlan 10 rx
Switch(config)# end
Switch# show monitor session 2
Session 2
---------
Source Ports:
    RX Only:       None
    TX Only:       None
    Both:          None
Source VLANs:
    RX Only:       1-3,10
    TX Only:       None
    Both:          None
Destination Ports:Gi0/7
    Encapsulation: Native
Filter VLANs:      None
```

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit SPAN source traffic to specific VLANs:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no monitor session** *session_number*	Clear any existing SPAN configuration for the session.
Step 3	**monitor session** *session_number* **interface** *interface-id* **rx**	Specify the characteristics of the source port (monitored port) and SPAN session. For *session_number*, specify 1 or 2. For *interface-id*, specify the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	**monitor session** *session_number* **filter vlan** *vlan-id* [**,** | **-**]	Limit the SPAN source traffic to specific VLANs. For *session_number*, specify 1 or 2. For *vlan-id*, the range is 1 to 1005. (Optional) Use a comma (**,**) to specify a series of VLANs or use a hyphen (**-**) to specify a range of VLANs. Enter a space after the comma; enter a space before and after the hyphen.
Step 5	**monitor session** *session_number* **destination interface** *interface-id*	Specify the characteristics of the destination port (monitoring port) and SPAN session. For *session_number*, specify 1 or 2. For *interface-id*, specify the destination port. Valid interfaces include physical interfaces.
Step 6	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 7	**show monitor** [**session** *session_number*]	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter** global configuration command.

This example shows how to clear any existing configuration on SPAN session 2, configure SPAN session 2 to monitor traffic received on trunk port 4, and send traffic for only VLANs 1 through 5 and 9 to destination port 8.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet0/4 rx
Switch(config)# monitor session 2 filter vlan 1 - 5 , 9
Switch(config)# monitor session 2 destination interface gigabitethernet0/8
Switch(config)# end
Switch# show monitor session 2
Session 2
---------
Source Ports:
    RX Only:       Gi0/4
    TX Only:       None
    Both:          None
Source VLANs:
    RX Only:       None
    TX Only:       None
    Both:          None
Destination Ports:Gi0/8
    Encapsulation: Native
Filter VLANs:      1-5,9
```

Displaying SPAN Status

To display the status of the current SPAN configuration, use the **show monitor** privileged EXEC command.

This is an example of output for the **show monitor** privileged EXEC command for **session 2**:

```
Switch# show monitor session 2
Session 2
---------
Source Ports:
    RX Only:       Gi0/4
    TX Only:       None
    Both:          None
Source VLANs:
    RX Only:       None
    TX Only:       None
    Both:          None
Destination Ports:Gi0/7
    Encapsulation: Native
Filter VLANs:      1-5,9
```

CHAPTER 16

Configuring RMON

This chapter describes how to configure Remote Network Monitoring (RMON) on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding RMON, page 16-1
- Configuring RMON, page 16-2
- Displaying RMON Status, page 16-6

Understanding RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON feature with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments.

Figure 16-1 Remote Monitoring Example

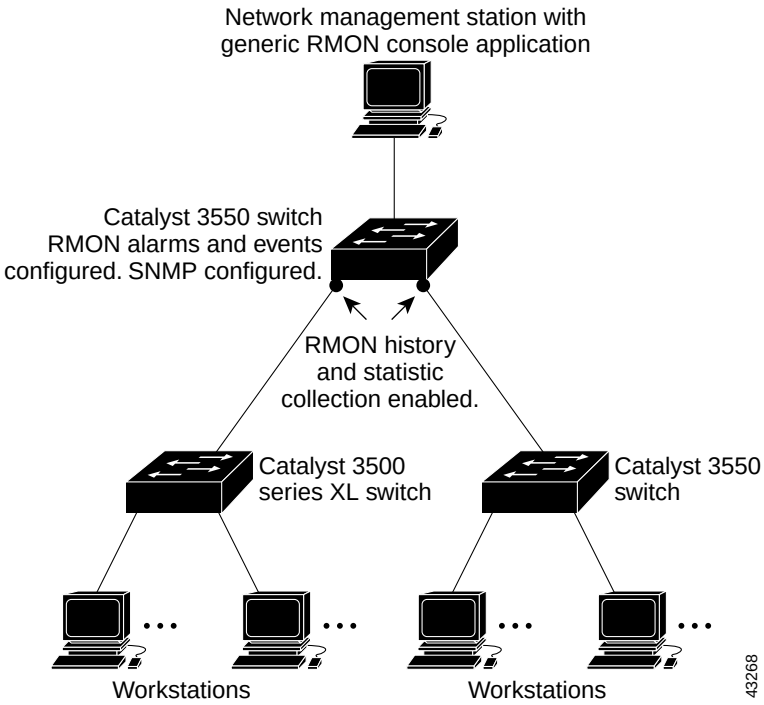

The switch supports these RMON groups (defined in RFC 1757):

- Statistics (RMON group 1)—Collects Ethernet, Fast Ethernet, and Gigabit Ethernet statistics on an interface.
- History (RMON group 2)—Collects a history group of statistics on Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces for a specified polling interval.
- Alarm (RMON group 3)—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Event (RMON group 9)—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry or an SNMP trap.

Because switches supported by this IOS release use hardware counters for RMON data processing, the monitoring is more efficient, and little processing power is required.

Configuring RMON

This section describes how to configure RMON on your switch. It contains this configuration information:

- Default RMON Configuration, page 16-3
- Configuring RMON Alarms and Events, page 16-3
- Configuring RMON Collection on an Interface, page 16-5

Default RMON Configuration

RMON is disabled by default; no alarms or events are configured.

Only RMON 1 is supported on the switch.

Configuring RMON Alarms and Events

You can configure your switch for RMON by using the command-line interface (CLI) or an SNMP-compatible network management station. We recommend that you use a generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. You must also configure SNMP on the switch to access RMON MIB objects. For more information, see Chapter 18, "Configuring SNMP."

Beginning in privileged EXEC mode, follow these steps to enable RMON alarms and events:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**rmon alarm** *number variable interval* {**delta** | **absolute**} **rising-threshold** *value* [*event-number*] **falling-threshold** *value* [*event-number*] [**owner** *string*]	Set an alarm on a MIB object. • For *number*, specify the alarm number. The range is 1 to 65535. • For *variable*, specify the MIB object to monitor. • For *interval*, specify the time in seconds the alarm monitors the MIB variable. The range is 1 to 4294967295 seconds. • Specify the **delta** keyword to test the change between samples of a MIB variable; specify the **absolute** keyword to test each MIB variable directly. • For *value*, specify a number at which the alarm is triggered and one for when the alarm is reset. The range for the rising threshold and falling threshold *values* is -2147483648 to 2147483647. • (Optional) For *event-number*, specify the event number to trigger when the rising or falling threshold exceeds its limit. • (Optional) For *string*, specify the owner of the alarm.

	Command	Purpose
Step 3	**rmon event** *number* [**log**] [**trap** *community*] [**description** *string*] [**owner** *string*]	Add an event in the RMON event table that is associated with an RMON event number. • For *number*, assign an event number. The range is 1 to 65535. • (Optional) Use the **log** keyword to generate an RMON log entry when the event is triggered. • (Optional) For *community*, enter the SNMP community string used for this trap. • (Optional) For **description** *string*, specify a description of the event. • (Optional) For **owner** *string*, specify the owner of this event.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable an alarm, use the **no rmon alarm** *number* global configuration command on each alarm you configured. You cannot disable at once all the alarms that you configured. To disable an event, use the **no rmon event** *number* global configuration command. To learn more about alarms and events and how they interact with each other, refer to RFC 1757.

You can set an alarm on any MIB object. The following example configures RMON alarm number 10 by using the **rmon alarm** command. The alarm monitors the MIB variable *ifEntry.20.1* once every 20 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the *ifEntry.20.1* value shows a MIB counter increase of 15 or more, such as from 100000 to 100015, the alarm is triggered. The alarm in turn triggers event number 1, which is configured with the **rmon event** command. Possible events can include a log entry or an SNMP trap. If the *ifEntry.20.1* value changes by 0, the alarm is reset and can be triggered again.

```
Switch(config)# rmon alarm 10 ifEntry.20.1 20 delta rising-threshold 15 1
falling-threshold 0 owner jjohnson
```

The following example creates RMON event number 1 by using the **rmon event** command. The event is defined as *High ifOutErrors* and generates a log entry when the event is triggered by the alarm. The user *jjones* owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.

```
Switch(config)# rmon event 1 log trap eventtrap description "High ifOutErrors" owner
jjones
```

Configuring RMON Collection on an Interface

You must first configure RMON alarms and events to display collection information.

Beginning in privileged EXEC mode, follow these steps to collect group history statistics on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface on which to collect history.
Step 3	**rmon collection history** *index* [**owner** *ownername*] [**buckets** *bucket-number*] [**interval** *seconds*]	Enable history collection for the specified number of buckets and time period. • For *index*, identify the RMON group of statistics The range is 1 to 65535. • (Optional) For *ownername*, enter the name of the owner of the RMON group of statistics. • For *bucket-number*, specify the maximum number of buckets desired for the RMON collection history group of statistics. The range is 1 to 65535. The default is 50 buckets. • For *seconds*, specify the number of seconds in each polling cycle.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**show rmon history**	Display the contents of the switch history table.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable history collection, use the **no rmon collection history** *index* interface configuration command.

Beginning in privileged EXEC mode, follow these steps to collect group Ethernet statistics on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface on which to collect statistics.
Step 3	**rmon collection stats** *index* [**owner** *ownername*]	Enable RMON statistic collection on the interface. • For *index*, specify the RMON group of statistics. The range is from 1 to 65535. • (Optional) For *ownername*, enter the name of the owner of the RMON group of statistics.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.

	Command	Purpose
Step 6	**show rmon statistics**	Display the contents of the switch statistics table.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the collection of group Ethernet statistics, use the **no rmon collection stats** *index* interface configuration command.

Displaying RMON Status

To display the RMON status, use one or more of the privileged EXEC commands in Table 16-1:

Table 16-1 Commands for Displaying RMON Status

Command	Purpose
show rmon	Displays general RMON statistics.
show rmon alarms	Displays the RMON alarm table.
show rmon events	Displays the RMON event table.
show rmon history	Displays the RMON history table.
show rmon statistics	Displays the RMON statistics table.

CHAPTER 17

Configuring System Message Logging

This chapter describes how to configure system message logging on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding System Message Logging, page 17-1
- Configuring System Message Logging, page 17-2
- Displaying the Logging Configuration, page 17-12

Understanding System Message Logging

By default, switches send the output from system messages and **debug** privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a UNIX syslog server, depending on your configuration. The process also sends messages to the console.

Note The syslog format is compatible with 4.3 BSD UNIX.

When the logging process is disabled, messages are sent only to the console. The messages are sent as they are generated, so message and debug output are interspersed with prompts or output from other commands. Messages are displayed on the console after the process that generated them has finished.

You can set the severity level of the messages to control the type of messages displayed on the console and each of the destinations. You can timestamp log messages or set the syslog source address to enhance real-time debugging and management. For information on possible messages, refer to the *Catalyst 3550 Multilayer Switch System Message Guide* for this release.

You can access logged system messages by using the switch command-line interface (CLI) or by saving them to a properly configured syslog server. The switch software saves syslog messages in an internal buffer. You can remotely monitor system messages by accessing the switch through Telnet, through the console port, or by viewing the logs on a syslog server.

Configuring System Message Logging

This section describes how to configure system message logging. It contains this configuration information:

- System Log Message Format, page 17-2
- Default System Message Logging Configuration, page 17-3
- Disabling and Enabling Message Logging, page 17-4
- Setting the Message Display Destination Device, page 17-4
- Synchronizing Log Messages, page 17-6
- Enabling and Disabling Timestamps on Log Messages, page 17-7
- Enabling and Disabling Sequence Numbers in Log Messages, page 17-8
- Defining the Message Severity Level, page 17-8
- Limiting Syslog Messages Sent to the History Table and to SNMP, page 17-10
- Configuring UNIX Syslog Servers, page 17-10

System Log Message Format

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or timestamp information, if configured. Messages are displayed in this format:

seq no:timestamp: %facility-severity-MNEMONIC:description

The part of the message preceding the percent sign depends on the setting of the **service sequence-numbers**, **service timestamps log datetime**, **service timestamps log datetime** [**localtime**] [**msec**] [**show-timezone**], or **service timestamps log uptime** global configuration command.

Table 17-1 describes the elements of syslog messages.

Table 17-1 System Log Message Elements

Element	Description
seq no:	Stamps log messages with a sequence number only if the **service sequence-numbers** global configuration command is configured. For more information, see the "Enabling and Disabling Sequence Numbers in Log Messages" section on page 17-8.
timestamp formats: *mm/dd hh:mm:ss* or *hh:mm:ss* (short uptime) or *d h* (long uptime)	Date and time of the message or event. This information appears only if the **service timestamps log** [**datetime** | **log**] global configuration command is configured. For more information, see the "Enabling and Disabling Timestamps on Log Messages" section on page 17-7.
facility	The facility to which the message refers (for example, SNMP, SYS, and so forth). For a list of supported facilities, see Table 17-4 on page 17-12.
severity	Single-digit code from 0 to 7 that is the severity of the message. For a description of the severity levels, see Table 17-3 on page 17-9.

Table 17-1 System Log Message Elements (continued)

Element	Description
MNEMONIC	Text string that uniquely describes the message.
description	Text string containing detailed information about the event being reported.

This example shows a partial switch system message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed
state to down 2
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
18:47:02: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
*Mar  1 18:48:50.483 UTC: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Default System Message Logging Configuration

Table 17-2 shows the default system message logging configuration.

Table 17-2 Default System Message Logging Configuration

Feature	Default Setting
System message logging to the console	Enabled.
Console severity	Debugging (and numerically lower levels; see Table 17-3 on page 17-9).
Logging buffer size	4096 bytes.
Logging history size	1 message.
Timestamps	Disabled.
Synchronous logging	Disabled.
Logging server	Disabled.
Syslog server IP address	None configured.
Server facility	Local7 (see Table 17-4 on page 17-12).
Server severity	Informational (and numerically lower levels; see Table 17-3 on page 17-9).

Disabling and Enabling Message Logging

Message logging is enabled by default. It must be enabled to send messages to any destination other than the console. When enabled, log messages are sent to a logging process, which logs messages to designated locations asynchronously to the processes that generated the messages.

Beginning in privileged EXEC mode, follow these steps to disable message logging:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no logging on**	Disable message logging.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config** or **show logging**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Disabling the logging process can slow down the switch because a process must wait until the messages are written to the console before continuing. When the logging process is disabled, messages are displayed on the console as soon as they are produced, often appearing in the middle of command output.

The **logging synchronous** global configuration command also affects the display of messages to the console. When this command is enabled, messages appear only after you press Return. For more information, see the "Synchronizing Log Messages" section on page 17-6.

To re-enable message logging after it has been disabled, use the **logging on** global configuration command.

Setting the Message Display Destination Device

If message logging is enabled, you can send messages to specific locations in addition to the console. Beginning in privileged EXEC mode, use one or more of the following commands to specify the locations that receive messages:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**logging buffered** [*size*]	Log messages to an internal buffer. The default buffer size is 4096. The range is 4096 to 4294967295 bytes. **Note** Do not make the buffer size too large because the switch could run out of memory for other tasks. Use the **show memory** privileged EXEC command to view the free processor memory on the switch; however, this value is the maximum available, and the buffer size should *not* be set to this amount.

	Command	Purpose
Step 3	**logging** *host*	Log messages to a UNIX syslog server host. For *host*, specify the name or IP address of the host to be used as the syslog server. To build a list of syslog servers that receive logging messages, enter this command more than once. For complete syslog server configuration steps, see the "Configuring UNIX Syslog Servers" section on page 17-10.
Step 4	**logging file flash:***filename* [*max-file-size*] [*min-file-size*] [*severity-level-number* | *type*]	Store log messages in a file in Flash memory. • For *filename*, enter the log message filename. • (Optional) For *max-file-size*, specify the maximum logging file size. The range is 4096 to 2147483647. The default is 4069 bytes. • (Optional) For *min-file-size*, specify the minimum logging file size. The range is 1024 to 2147483647. The default is 2048 bytes. • (Optional) For *severity-level-number* | *type*, specify either the logging severity level or the logging type. The severity range is 0 to 7. For a list of logging type keywords, see Table 17-3 on page 17-9. By default, the log file receives debugging messages and numerically lower levels.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**terminal monitor**	Log messages to a nonconsole terminal during the current session. Terminal parameter-setting commands are set locally and do not remain in effect after the session has ended. You must perform this step for each session to see the debugging messages.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

The **logging buffered** global configuration command copies logging messages to an internal buffer. The buffer is circular, so newer messages overwrite older messages after the buffer is full. To display the messages that are logged in the buffer, use the **show logging** privileged EXEC command. The first message displayed is the oldest message in the buffer. To clear the contents of the buffer, use the **clear logging** privileged EXEC command.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a file, use the **no logging file** [*severity-level-number* | *type*] global configuration command.

Synchronizing Log Messages

You can configure the system to synchronize unsolicited messages and **debug** privileged EXEC command output with solicited device output and prompts for a specific console port line or virtual terminal line. You can identify the types of messages to be output asynchronously based on the level of severity. You can also determine the maximum number of buffers for storing asynchronous messages for the terminal after which messages are dropped.

When synchronous logging of unsolicited messages and **debug** command output is enabled, unsolicited device output is displayed on the console or printed after solicited device output is displayed or printed. Unsolicited messages and **debug** command output is displayed on the console after the prompt for user input is returned. Therefore, unsolicited messages and **debug** command output are not interspersed with solicited device output and prompts. After the unsolicited messages are displayed, the console again displays the user prompt.

Beginning in privileged EXEC mode, follow these steps to configure synchronous logging:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**line** [**console** | **vty**] *line-number* [*ending-line-number*]	Specify the line to be configured for synchronous logging of messages. • Use the **console** keyword for configurations that occur through the switch console port. • Use the **line vty** *line-number* command to specify which vty lines are to have synchronous logging enabled. You use a vty connection for configurations that occur through a Telnet session. The range of line numbers is from 0 to 15. You can change the setting of all 16 vty lines at once by entering: **line vty 0 15** Or you can change the setting of the single vty line being used for your current connection. For example, to change the setting for vty line 2, enter: **line vty 2** Entering this command changes to line configuration mode.
Step 3	**logging synchronous** [**level** *severity-level* | **all**] [**limit** *number-of-buffers*]	Enable synchronous logging of messages. • (Optional) For **level** *severity-level*, specify the message severity level. Messages with a severity level equal to or higher than this value are printed asynchronously. Low numbers mean greater severity and high numbers mean lesser severity. The default is 2. • (Optional) Specifying **level all** means that all messages are printed asynchronously regardless of the severity level. • (Optional) For **limit** *number-of-buffers*, specify the number of buffers to be queued for the terminal after which new messages are dropped. The default is 20.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable synchronization of unsolicited messages and debug output, use the **no logging synchronous** [**level** *severity-level* | **all**] [**limit** *number-of-buffers*] line configuration command.

Enabling and Disabling Timestamps on Log Messages

By default, log messages are not timestamped.

Beginning in privileged EXEC mode, follow these steps to enable timestamping of log messages:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**service timestamps log uptime** or **service timestamps log datetime** [**msec**] [**localtime**] [**show-timezone**]	Enable log timestamps. The first command enables timestamps on log messages, showing the time since the system was rebooted. The second command enables timestamps on log messages. Depending on the options selected, the timestamp can include the date, time in milliseconds relative to the local time zone, and the time zone name.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable timestamps for both debug and log messages, use the **no service timestamps** global configuration command.

This example shows part of a logging display with the **service timestamps log datetime** global configuration command enabled:

```
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

This example shows part of a logging display with the s**ervice timestamps log uptime** global configuration command enabled:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
```

Enabling and Disabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same timestamp, you can display messages with sequence numbers so that you can unambiguously refer to a single message. By default, sequence numbers in log messages are not displayed.

Beginning in privileged EXEC mode, follow these steps to enable sequence numbers in log messages:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**service sequence-numbers**	Enable sequence numbers.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable sequence numbers, use the **no service sequence-numbers** global configuration command.

This example shows part of a logging display with sequence numbers enabled:

```
000019: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Defining the Message Severity Level

You can limit messages displayed to the selected device by specifying the severity level of the message, which are described in Table 17-3.

Beginning in privileged EXEC mode, follow these steps to define the message severity level:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**logging console** *level*	Limit messages logged to the console. By default, the console receives debugging messages and numerically lower levels (see Table 17-3 on page 17-9).
Step 3	**logging monitor** *level*	Limit messages logged to the terminal lines. By default, the terminal receives debugging messages and numerically lower levels (see Table 17-3 on page 17-9).
Step 4	**logging trap** *level*	Limit messages logged to the syslog servers. By default, syslog servers receive informational messages and numerically lower levels (see Table 17-3 on page 17-9). For complete syslog server configuration steps, see the "Configuring UNIX Syslog Servers" section on page 17-10.
Step 5	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 6	**show running-config** or **show logging**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note Specifying a *level* causes messages at that level and numerically lower levels to be displayed at the destination.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a terminal other than the console, use the **no logging monitor** global configuration command. To disable logging to syslog servers, use the **no logging trap** global configuration command.

Table 17-3 describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 17-3 Message Logging Level Keywords

Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

The software generates four other categories of messages:

- Error messages about software or hardware malfunctions, displayed at levels **warnings** through **emergencies**. These types of messages mean that the functionality of the switch is affected. For information on how to recover from these malfunctions, refer to the *Catalyst 3550 Multilayer Switch System Message Guide*.
- Output from the **debug** commands, displayed at the **debugging** level. Debug commands are typically used only by the Technical Assistance Center.
- Interface up or down transitions and system restart messages, displayed at the **notifications** level. This message is only for information; switch functionality is not affected.
- Reload requests and low-process stack messages, displayed at the **informational** level. This message is only for information; switch functionality is not affected.

Limiting Syslog Messages Sent to the History Table and to SNMP

If you have enabled syslog message traps to be sent to an SNMP network management station by using the **snmp-server enable trap** global configuration command, you can change the level of messages sent and stored in the switch history table. You can also change the number of messages that are stored in the history table.

Messages are stored in the history table because SNMP traps are not guaranteed to reach their destination. By default, one message of the level **warning** and numerically lower levels (see Table 17-3 on page 17-9) are stored in the history table even if syslog traps are not enabled.

Beginning in privileged EXEC mode, follow these steps to change the level and history table size defaults:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**logging history** *level*[1]	Change the default level of syslog messages stored in the history file and sent to the SNMP server. See Table 17-3 on page 17-9 for a list of *level* keywords. By default, **warnings**, **errors**, **critical**, **alerts**, and **emergencies** messages are sent.
Step 3	**logging history size** *number*	Specify the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 1 to 500 messages.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

1. Table 17-3 lists the level keywords and severity level. For SNMP usage, the severity level values increase by 1. For example, emergencies equal 1, not 0, and critical equals 3, not 2.

When the history table is full (it contains the maximum number of message entries specified with the **logging history size** global configuration command), the oldest message entry is deleted from the table to allow the new message entry to be stored.

To return the logging of syslog messages to the default level, use the **no logging history** global configuration command. To return the number of messages in the history table to the default value, use the **no logging history size** global configuration command.

Configuring UNIX Syslog Servers

The next sections describe how to configure the UNIX server syslog daemon and define the UNIX system logging facility.

Logging Messages to a UNIX Syslog Daemon

Before you can send system log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. Log in as root, and perform these steps:

Note Some recent versions of UNIX syslog daemons no longer accept by default syslog packets from the network. If this is the case with your system, use the UNIX **man syslogd** command to determine what options must be added to or removed from the syslog command line to enable logging of remote syslog messages.

Step 1 Add a line such as the following to the file /etc/syslog.conf:

```
local7.debug /usr/adm/logs/cisco.log
```

The **local7** keyword specifies the logging facility to be used; see Table 17-4 on page 17-12 for information on the facilities. The **debug** keyword specifies the syslog level; see Table 17-3 on page 17-9 for information on the severity levels. The syslog daemon sends messages at this level or at a more severe level to the file specified in the next field. The file must already exist, and the syslog daemon must have permission to write to it.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/cisco.log
$ chmod 666 /var/log/cisco.log
```

Step 3 Make sure the syslog daemon reads the new changes by entering this command:

```
$ kill -HUP `cat /etc/syslog.pid`
```

For more information, see the **man syslog.conf** and **man syslogd** commands on your UNIX system.

Configuring the UNIX System Logging Facility

When sending system log messages to an external device, you can cause the switch to identify its messages as originating from any of the UNIX syslog facilities.

Beginning in privileged EXEC mode, follow these steps to configure UNIX system facility message logging:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**logging** *host*	Log messages to a UNIX syslog server host by entering its IP address. To build a list of syslog servers that receive logging messages, enter this command more than once.
Step 3	**logging trap** *level*	Limit messages logged to the syslog servers. Be default, syslog servers receive informational messages and lower. See Table 17-3 on page 17-9 for *level* keywords.

	Command	Purpose
Step 4	**logging facility** *facility-type*	Configure the syslog facility. See Table 17-4 on page 17-12 for *facility-type* keywords. The default is **local7**.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a syslog server, use the **no logging** *host* global configuration command, and specify the syslog server IP address. To disable logging to syslog servers, enter the **no logging trap** global configuration command.

Table 17-4 lists the UNIX system facilities supported by the Cisco IOS software. For more information about these facilities, consult the operator's manual for your UNIX operating system.

Table 17-4 Logging Facility-Type Keywords

Facility Type Keyword	Description
auth	Authorization system
cron	Cron facility
daemon	System daemon
kern	Kernel
local0-7	Locally defined messages
lpr	Line printer system
mail	Mail system
news	USENET news
sys9	System use
sys10	System use
sys11	System use
sys12	System use
sys13	System use
sys14	System use
syslog	System log
user	User process
uucp	UNIX-to-UNIX copy system

Displaying the Logging Configuration

To display the current logging configuration and the contents of the log buffer, use the **show logging** privileged EXEC command. For information about the fields in this display, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

CHAPTER 18

Configuring SNMP

This chapter describes how to configure the Simple Network Management Protocol (SNMP) on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding SNMP, page 18-1
- Configuring SNMP, page 18-4
- Displaying SNMP Status, page 18-10

Understanding SNMP

SNMP is an application-layer protocol that provides a message format for communication between SNMP managers and agents. The SNMP manager can be part of a network management system (NMS) such as CiscoWorks. The agent and management information base (MIB) reside on the switch. To configure SNMP on the switch, you define the relationship between the manager and the agent.

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into the agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can mean improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a TCP connection, loss of connection to a neighbor, or other significant events.

This section includes this conceptual information:

- SNMP Versions, page 18-2
- SNMP Manager Functions, page 18-2
- SNMP Agent Functions, page 18-3
- SNMP Community Strings, page 18-3
- Using SNMP to Access MIB Variables, page 18-3

SNMP Versions

This software release supports these SNMP versions:

- SNMPv1—The Simple Network Management Protocol, a Full Internet Standard, defined in RFC 1157.
- SNMPv2C, which has these features:
  - SNMPv2—Version 2 of the Simple Network Management Protocol, a Draft Internet Standard, defined in RFCs 1902 through 1907.
  - SNMPv2C—The Community-based Administrative Framework for SNMPv2, an Experimental Internet Protocol defined in RFC 1901.

SNMPv2C replaces the Party-based Administrative and Security Framework of SNMPv2Classic with the Community-based Administrative Framework of SNMPv2C while retaining the bulk retrieval and improved error handling of SNMPv2Classic.

Both SNMPv1 and SNMPv2C use a community-based form of security. The community of managers able to access the agent's MIB is defined by an IP address access control list and password.

SNMPv2C includes a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism retrieves tables and large quantities of information, minimizing the number of round-trips required. The SNMPv2C improved error-handling includes expanded error codes that distinguish different kinds of error conditions; these conditions are reported through a single error code in SNMPv1. Error return codes now report the error type.

You must configure the SNMP agent to use the version of SNMP supported by the management station. An agent can communicate with multiple managers; for this reason, you can configure the software to support communications with one management station using the SNMPv1 protocol and another using the SNMPv2 protocol.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 18-1.

Table 18-1 SNMP Operations

Operation	Description
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table.[1]
get-bulk-request[2]	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data.
get-response	Replies to a get-request, get-next-request, and set-request sent by an NMS.
set-request	Stores a value in a specific variable.
trap	An unsolicited message sent by an SNMP agent to an SNMP manager when some event has occurred.

1. With this operation, an SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.
2. The **get-bulk** command only works with SNMPv2.

SNMP Agent Functions

The SNMP agent responds to SNMP manager requests as follows:

- Get a MIB variable—The SNMP agent begins this function in response to a request from the NMS. The agent retrieves the value of the requested MIB variable and responds to the NMS with that value.
- Set a MIB variable—The SNMP agent begins this function in response to a message from the NMS. The SNMP agent changes the value of the MIB variable to the value requested by the NMS.

The SNMP agent also sends unsolicited trap messages to notify an NMS that a significant event has occurred on the agent. Examples of trap conditions include, but are not limited to, when a port or module goes up or down, when spanning-tree topology changes occur, and when authentication failures occur.

SNMP Community Strings

SNMP community strings authenticate access to MIB objects and function as embedded passwords. In order for the NMS to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the switch.

A community string can have one of these attributes:

- Read-only—Gives read access to authorized management stations to all objects in the MIB except the community strings, but does not allow write access
- Read-write—Gives read and write access to authorized management stations to all objects in the MIB, but does not allow access to the community strings
- Read-write-all—Gives read and write access to authorized management stations to all objects in the MIB, including the community strings

Note When a cluster is created, the command switch manages the exchange of messages among member switches and the SNMP application. The Cluster Management software appends the member switch number (*@esN*, where *N* is the switch number) to the first configured RW and RO community strings on the command switch and propagates them to the member switches. For more information, see Chapter 5, “Clustering Switches.”

Using SNMP to Access MIB Variables

An example of an NMS is the CiscoWorks network management software. CiscoWorks 2000 software uses the switch MIB variables to set device variables and to poll devices on the network for specific information. The results of a poll can be displayed as a graph and analyzed to troubleshoot internetworking problems, increase network performance, verify the configuration of devices, monitor traffic loads, and more.

As shown in Figure 18-1, the SNMP agent gathers data from the MIB. The agent can send traps, or notification of certain events, to the SNMP manager, which receives and processes the traps. Traps are messages alerting the SNMP manager to a condition on the network such as improper user authentication, restarts, link status (up or down), MAC address tracking, and so forth. The SNMP agent also responds to MIB-related queries sent by the SNMP manager in *get-request*, *get-next-request*, and *set-request* format.

Figure 18-1 SNMP Network

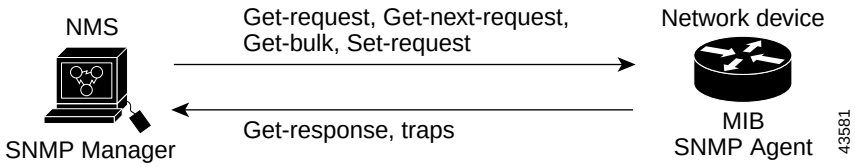

For information on supported MIBs and how to access them, see Appendix A, "Supported MIBs."

Configuring SNMP

This section describes how to configure SNMP on your switch. It contains this configuration information:

- Default SNMP Configuration, page 18-4
- Configuring Community Strings, page 18-5
- Configuring Trap Managers and Enabling Traps, page 18-7
- Setting the Agent Contact and Location Information, page 18-9
- Limiting TFTP Servers Used Through SNMP, page 18-9
- SNMP Examples, page 18-10

Default SNMP Configuration

Table 18-2 shows the default SNMP configuration.

Table 18-2 Default SNMP Configuration

Feature	Default Setting
SNMP agent	Enabled
SNMP community strings	Read-Only: Public Read-Write: Private Read-Write-all: Secret
SNMP trap receiver	None configured
SNMP traps	None enabled

Disabling the SNMP Agent

Beginning in privileged EXEC mode, follow these steps to disable the SNMP agent:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no snmp-server**	Disable the SNMP agent operation.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

No specific IOS command exists to enable SNMP. The first **snmp-server** global configuration command that you enter enables SNMPv1 and SNMPv2.

Configuring Community Strings

You use the SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the switch. Optionally, you can specify one or more of these characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent
- A MIB view, which defines the subset of all MIB objects accessible to the given community
- Read and write or read-only permission for the MIB objects accessible to the community

Beginning in privileged EXEC mode, follow these steps to configure a community string on the switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**snmp-server community** *string* [**ro** | **rw**] [*access-list-number*]	Configure the community string. • For *string*, specify a string that acts like a password and permits access to the SNMP protocol. You can configure one or more community strings of any length. • (Optional) Specify either read-only (**ro**) if you want authorized management stations to retrieve MIB objects, or specify read-write (**rw**) if you want authorized management stations to retrieve and modify MIB objects. By default, the community string permits read-only access to all objects. • (Optional) For *access-list-number*, enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.

	Command	Purpose
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	(Optional) If you specified an IP standard access list number in Step 2, then create the list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the IP address of the SNMP managers that are permitted to use the community string to gain access to the agent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note To disable access for an SNMP community, set the community string for that community to the null string (do not enter a value for the community string).

To remove a specific community string, use the **no snmp-server community** *string* global configuration command.

This example shows how to assign the string *comaccess* to SNMP, to allow read-only access, and to specify that IP access list 4 can use the community string to gain access to the switch SNMP agent:

```
Switch(config)# snmp-server community comaccess ro 4
```

Configuring Trap Managers and Enabling Traps

A trap manager is a management station that receives and processes traps. Traps are system alerts that the switch generates when certain events occur. By default, no trap manager is defined, and no traps are issued.

Switches running this IOS release can have an unlimited number of trap managers. Community strings can be any length.

Table 18-3 describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them.

Table 18-3 Switch Notification Types

Notification Type	Description
cluster	Generates a trap when the cluster configuration changes.
config	Generates a trap for SNMP configuration changes.
entity	Generates a trap for SNMP entity changes.
HSRP	Generates a trap for Hot Standby Router Protocol (HSRP) changes.
MAC notification	Generates a trap for MAC address notifications.
RTR	Generates a trap for the SNMP Response Time Reporter (RTR).
SNMP	Generates a trap for SNMP-type notifications.
TTY	Sends Cisco enterprise-specific notifications when a Transmission Control Protocol (TCP) connection closes.
UDP-port	Sends notification of the User Datagram Protocol (UDP) port number of the host.
VTP	Generates a trap for VLAN Trunking Protocol (VTP) changes.

Note Though visible in the command-line help string, the **hsrp** keyword takes affect only when the enhanced multilayer software image (EMI) is installed. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Some notification types cannot be controlled with the **snmp-server enable** global configuration command, for example, **tty** and **udp-port**. These notification types are always enabled. You can use the **snmp-server host** global configuration command to a specific host to receive the notification types listed in Table 18-3.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send traps to a host:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**snmp-server host** *host-addr* {**traps** | **informs**} {**version** {**1** | **2c**}} *community-string notification-type*	Specify the recipient of the trap message. • For *host-addr*, specify the name or address of the host (the targeted recipient). • Specify **traps** (the default) to send SNMP traps to the host. Specify **informs** to send SNMP informs to the host. • Specify the SNMP version to support. Version 1, the default, is not available with informs. **Note** Though visible in the command-line help string, the **version 3** keyword (SNMPv3) is not supported. • For *community-string*, specify the string to send with the notification operation. Though you can set this string using the **snmp-server host** command, we recommend that you define this string by using the **snmp-server community** command before using the **snmp-server host** command. • For *notification-type*, use the keywords listed in Table 18-3 on page 18-7.
Step 3	**snmp-server enable traps** *notification-types*	Enable the switch to send specific traps. For a list of traps, see Table 18-3 on page 18-7. To enable multiple types of traps, you must issue a separate **snmp-server enable traps** command for each trap type.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the specified host from receiving traps, use the **no snmp-server host** *host* global configuration command. To disable a specific trap type, use the **no snmp-server enable traps** *notification-types* global configuration command.

Setting the Agent Contact and Location Information

Beginning in privileged EXEC mode, follow these steps to set the system contact and location of the SNMP agent so that these descriptions can be accessed through the configuration file:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**snmp-server contact** *text*	Set the system contact string. For example: `snmp-server contact Dial System Operator at beeper 21555.`
Step 3	**snmp-server location** *text*	Set the system location string. For example: `snmp-server location Building 3/Room 222`
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Limiting TFTP Servers Used Through SNMP

Beginning in privileged EXEC mode, follow these steps to limit the TFTP servers used for saving and loading configuration files through SNMP to the servers specified in an access list:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**snmp-server tftp-server-list** *access-list-number*	Limit TFTP servers used for configuration file copies through SNMP to the servers in the access list. For *access-list-number*, enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the IP address of the TFTP servers that can access the switch. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.

	Command	Purpose
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

SNMP Examples

This example shows how to enable SNMPv1 and SNMPv2C. The configuration permits any SNMP manager to access all objects with read-only permissions using the community string *public*. This configuration does not cause the switch to send any traps.

```
Switch(config)# snmp-server community public
```

This example shows how to permit any SNMP manager to access all objects with read-only permission using the community string *public*. The switch also sends VTP traps to the hosts 192.180.1.111 and 192.180.1.33 using SNMPv1 and to the host 192.180.1.27 using SNMPv2C. The community string *public* is sent with the traps.

```
Switch(config)# snmp-server community public
Switch(config)# snmp-server enable traps vtp
Switch(config)# snmp-server host 192.180.1.27 version 2c public
Switch(config)# snmp-server host 192.180.1.111 version 1 public
Switch(config)# snmp-server host 192.180.1.33 public
```

This example shows how to allow read-only access for all objects to members of access list 4 that use the *comaccess* community string. No other SNMP managers have access to any objects. SNMP Authentication Failure traps are sent by SNMPv2C to the host *cisco.com* using the community string *public*.

```
Switch(config)# snmp-server community comaccess ro 4
Switch(config)# snmp-server enable traps snmp authentication
Switch(config)# snmp-server host cisco.com version 2c public
```

This example shows how to send Entity MIB traps to the host *cisco.com*. The community string is restricted. The first line enables the switch to send Entity MIB traps in addition to any traps previously enabled. The second line specifies the destination of these traps and overwrites any previous **snmp-server host** commands for the host *cisco.com*.

```
Switch(config)# snmp-server enable traps entity
Switch(config)# snmp-server host cisco.com restricted entity
```

This example shows how to enable the switch to send all traps to the host *myhost.cisco.com* using the community string *public*:

```
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server host myhost.cisco.com public
```

Displaying SNMP Status

To display SNMP input and output statistics, including the number of illegal community string entries, errors, and requested variables, use the **show snmp** privileged EXEC command. For information about the fields in this display, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

CHAPTER 19

Configuring Network Security with ACLs

This chapter describes how to configure network security on your switch by using access control lists (ACLs), which are also referred to in commands and tables as access lists. To take advantage of some of the features described in this chapter, you must have the enhanced multilayer software image installed on your switch.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the "Configuring IP Services" section of the *Cisco IOS IP and IP Routing Configuration Guide* and the *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.

This chapter consists of these sections:

- Understanding ACLs, page 19-1
- Configuring Router ACLs, page 19-5
- Configuring VLAN Maps, page 19-27
- Using VLAN Maps with Router ACLs, page 19-36

Note To allocate system resources to maximize the number of security access control entries (ACEs) allowed on the switch, you can use the **sdm prefer access** global configuration command to set the Switch Database Management feature to the access template. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

Understanding ACLs

Packet filtering can help limit network traffic and restrict network use by certain users or devices. ACLs can filter traffic as it passes through a router and permit or deny packets from crossing specified interfaces. An ACL is a sequential collection of permit and deny conditions that apply to packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. It tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packets. Because the switch stops testing conditions after the first match, the order of conditions in the list is critical. If no conditions match, the switch rejects the packets. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet.

Switches traditionally operate at Layer 2 only, switching traffic within a VLAN, whereas routers route traffic between VLANs. The Catalyst 3550 switch with the enhanced multilayer software image installed can accelerate packet routing between VLANs by using Layer 3 switching. The switch bridges the packet, the packet is then routed internally without going to an external router, and then the packet is bridged again to send it to its destination. During this process, the switch can access-control all packets it switches, including packets bridged within a VLAN.

You configure access lists on a router or Layer 3 switch to provide basic security for your network. If you do not configure ACLs, all packets passing through the switch could be allowed onto all parts of the network. You can use ACLs to control which hosts can access different parts of a network or to decide which types of traffic are forwarded or blocked at router interfaces. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic. ACLs can be configured to block inbound traffic, outbound traffic, or both.

An ACL contains an ordered list of access control entries (ACEs). Each ACE specifies *permit* or *deny* and a set of conditions the packet must satisfy in order to match the ACE. The meaning of *permit* or *deny* depends on the context in which the ACL is used.

The switch supports two types of ACLs:

- IP ACLs filter IP traffic, including TCP, User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP).
- Ethernet ACLs filter non-IP traffic.

Supported ACLs

The switch supports two applications of ACLs to filter traffic:

- Router ACLs access-control routed traffic between VLANs. All Catalyst 3550 switches can create router ACLs, but you must have the enhanced multilayer software image on your switch to filter packets routed between VLANs.
- VLAN ACLs or VLAN maps access-control all packets (bridged and routed). You can use VLAN maps to filter traffic between devices in the same VLAN. You do not need the enhanced image to create or apply VLAN maps. VLAN maps are configured to provide access-control based on Layer 3 addresses for IP. Unsupported protocols are access-controlled through MAC addresses using Ethernet ACEs.

  After a VLAN map is applied to a VLAN, all packets (routed or bridged) entering the VLAN are checked against the VLAN map. Packets can either enter the VLAN through a switch port or through a routed port after being routed.

This switch also supports Quality of Service (QoS) classification ACLs. For more information, see the "Classification Based on QoS ACLs" section on page 20-7.

Router ACLs

You can apply router ACLs on switch virtual interfaces (SVIs), which are Layer 3 interfaces to VLANs; on physical Layer 3 interfaces; and on Layer 3 EtherChannel interfaces. Router ACLs are applied on interfaces for specific directions (inbound or outbound).

One ACL can be used with multiple features for a given interface, and one feature can use multiple ACLs. When a single router ACL is used by multiple features, it is examined multiple times.

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses and optional protocol type information for matching operations.

The switch examines ACLs associated with features configured on a given interface and a direction. As packets enter the switch on an interface, ACLs associated with all inbound features configured on that interface are examined. After packets are routed and before they are forwarded to the next hop, all ACLs associated with outbound features configured on the egress interface are examined.

ACLs permit or deny packet forwarding based on how the packet matches the entries in the ACL. For example, you can use access lists to allow one host to access a part of a network, but prevent another host from accessing the same part. In Figure 19-1, ACLs applied at the router input allow Host A to access the Human Resources network, but prevent Host B from accessing the same network.

Figure 19-1 Using ACLs to Control Traffic to a Network

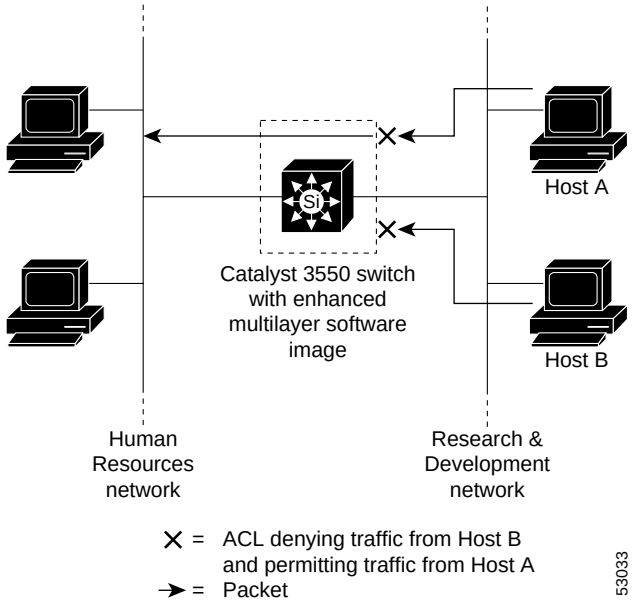

VLAN Maps

VLAN maps can access-control *all* traffic. You can apply VLAN maps on the switch to all packets that are routed into or out of a VLAN or are bridged within a VLAN. VLAN maps are used strictly for security packet filtering. Unlike router ACLs, VLAN maps are not defined by direction (input or output).

You can configure VLAN maps to match Layer 3 addresses for IP traffic. All non-IP protocols are access-controlled through MAC addresses and Ethertype using MAC VLAN maps. (IP traffic *is not* access controlled by MAC VLAN maps.) You can enforce VLAN maps only on packets going through the switch; you cannot enforce VLAN maps on traffic between hosts on a hub or on another switch connected to this switch.

With VLAN maps, forwarding of packets is permitted or denied, based on the action specified in the map. Figure 19-2 illustrates how a VLAN map is applied to deny a specific type of traffic from Host A in VLAN 10 from being forwarded.

Figure 19-2 Using VLAN Maps to Control Traffic

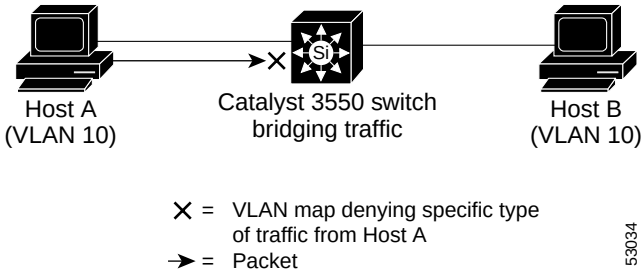

Handling Fragmented and Unfragmented Traffic

IP packets can be fragmented as they cross the network. When this happens, only the fragment containing the beginning of the packet contains the Layer 4 information, such as TCP or UDP port numbers, ICMP type and code, and so on. All other fragments are missing this information.

Some ACEs do not check Layer 4 information and therefore can be applied to all packet fragments. ACEs that do test Layer 4 information cannot be applied in the standard manner to most of the fragments in a fragmented IP packet. When the fragment contains no Layer 4 information and the ACE tests some Layer 4 information, the matching rules are modified:

- Permit ACEs that check the Layer 3 information in the fragment (including protocol type, such as TCP, UDP, and so on) are considered to match the fragment regardless of what the missing Layer 4 information might have been.
- Deny ACEs that check Layer 4 information never match a fragment unless the fragment contains Layer 4 information.

Consider access list 102, configured with these commands, applied to three fragmented packets:

```
Switch (config)# access-list 102 permit tcp any host 10.1.1.1 eq smtp
Switch (config)# access-list 102 deny tcp any host 10.1.1.2 eq telnet
Switch (config)# access-list 102 permit tcp any host 10.1.1.2
Switch (config)# access-list 102 deny tcp any any
```

Note In the first and second ACEs in the examples, the *eq* keyword after the destination address means to test for the TCP-destination-port well-known numbers equaling Simple Mail Transfer Protocol (SMTP) and Telnet, respectively.

- Packet A is a TCP packet from host 10.2.2.2., port 65000, going to host 10.1.1.1 on the SMTP port. If this packet is fragmented, the first fragment matches the first ACE (a permit) as if it were a complete packet because all Layer 4 information is present. The remaining fragments also match the first ACE, even though they do not contain the SMTP port information, because the first ACE only checks Layer 3 information when applied to fragments. The information in this example is that the packet is TCP and that the destination is 10.1.1.1.

- Packet B is from host 10.2.2.2, port 65001, going to host 10.1.1.2 on the Telnet port. If this packet is fragmented, the first fragment matches the second ACE (a deny) because all Layer 3 and Layer 4 information is present. The remaining fragments in the packet do not match the second ACE because they are missing Layer 4 information. Instead, they match the third ACE (a permit).

  Because the first fragment was denied, host 10.1.1.2 cannot reassemble a complete packet, so packet B is effectively denied. However, the later fragments that are permitted will consume bandwidth on the network and resources of host 10.1.1.2 as it tries to reassemble the packet.

- Fragmented packet C is from host 10.2.2.2, port 65001, going to host 10.1.1.3, port ftp. If this packet is fragmented, the first fragment matches the fourth ACE (a deny). All other fragments also match the fourth ACE because that ACE does not check any Layer 4 information and because Layer 3 information in all fragments shows that they are being sent to host 10.1.1.3, and the earlier permit ACEs were checking different hosts.

Configuring Router ACLs

Configuring router ACLs on Layer 3 switch-routed VLAN interfaces is the same as configuring ACLs on other Cisco routers. The process is briefly described here. For more detailed information on configuring router ACLs, refer to the "Configuring IP Services" chapter in the *Cisco IP and IP Routing Configuration Guide for IOS Release 12.1*. For detailed information about the commands, refer to *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*. For a list of IOS features not supported on the Catalyst 3550 switch, see the "Unsupported Features" section on page 19-6.

Caution By default, the router sends Internet Control Message Protocol (ICMP) unreachable messages when a packet is denied by an access group; these access-group denied packets are not dropped in hardware but are bridged to the switch CPU so that it can generate the ICMP-unreachable message. To drop access-group denied packets in hardware, you must disable ICMP unreachables by using the **no ip unreachables** interface configuration command. Note that the **ip unreachables** command is enabled by default.

This section includes the following information:

- Hardware and Software Handling of Router ACLs, page 19-5
- Unsupported Features, page 19-6
- Creating Standard and Extended IP ACLs, page 19-6
- Applying the ACL to an Interface or Terminal Line, page 19-18
- Displaying ACLs and Access Groups, page 19-20
- ACL Configuration Examples, page 19-22

Hardware and Software Handling of Router ACLs

ACL processing is primarily accomplished in hardware, but requires forwarding of some traffic flows to the CPU for software processing. The forwarding rate for software-forwarded traffic is substantially less than for hardware-forwarded traffic. When traffic flows are both logged and forwarded, forwarding is done by hardware, but logging must be done by software. Because of the difference in packet handling capacity between hardware and software, if the sum of all flows being logged (both permitted flows and denied flows) is of great enough bandwidth, not all of the packets that are forwarded can be logged.

These factors can cause packets to be sent to the CPU:

- Using the **log** keyword
- Enabling ICMP unreachables
- Hardware reaching its capacity to store ACL configurations

If ACLs cause large numbers of packets to be sent to the CPU, the switch performance can be negatively affected.

When you enter the **show ip access-lists** privileged EXEC command, the match count displayed does not account for packets that are access controlled in hardware. Use the **show access-lists hardware counters** privileged EXEC command to obtain some basic hardware ACL statistics for switched and routed packets.

Router ACLs function as follows:

- The hardware controls permit and deny actions of standard and extended ACLs (input and output) for security access control.
- If **log** has not been specified, the flows that match a *deny* statement in a security ACL are dropped by the hardware if *ip unreachables* is disabled. The flows matching a *permit* statement are switched in hardware.
- Adding the **log** keyword to an ACE in a router ACL causes a copy of the packet to be sent to the CPU only for logging. If the ACE is a *permit* statement, the packet is still switched and routed in hardware.

Unsupported Features

The Catalyst 3550 switch does not support these IOS router ACL-related features:

- Non-IP protocol ACLs (see Table 19-1 on page 19-7).
- Bridge-group ACLs.
- IP accounting.
- Inbound and outbound rate limiting (except with QoS ACLs).
- IP packets with a header length of less than five are not access controlled (results in an ICMP parameter error).
- Reflexive ACLs.
- Dynamic ACLs (except for certain specialized dynamic ACLs used by the switch clustering feature).

Creating Standard and Extended IP ACLs

This section summarizes how to create router IP ACLs. An ACL is a sequential collection of permit and deny conditions. The switch tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the switch denies the packet.

These are the steps to use ACLs:

Step 1 Create an ACL by specifying an access list number or name and access conditions.

Step 2 Apply the ACL to interfaces or terminal lines. You can also apply standard and extended IP ACLs to VLAN maps.

The software supports these styles of ACLs or access lists for IP:

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses for matching operations and optional protocol-type information for finer granularity of control.

These sections describe access lists and the steps for using them:

- Access List Numbers, page 19-7
- Creating a Numbered Standard ACL, page 19-8
- Creating a Numbered Extended ACL, page 19-9
- Creating Named Standard and Extended ACLs, page 19-14
- Applying Time Ranges to ACLs, page 19-15
- Including Comments About Entries in ACLs, page 19-18

Access List Numbers

The number you use to denote your ACL shows the type of access list that you are creating. Table 19-1 lists the access-list number and corresponding access list type and shows whether or not they are supported in the switch. The Catalyst 3550 switch supports IP standard and IP extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 19-1 Access List Numbers

Access List Number	Type	Supported
1–99	IP standard access list	Yes
100–199	IP extended access list	Yes
200–299	Protocol type-code access list	No
300–399	DECnet access list	No
400–499	XNS standard access list	No
500–599	XNS extended access list	No
600–699	AppleTalk access list	No
700–799	48-bit MAC address access list	No
800–899	IPX standard access list	No
900–999	IPX extended access list	No
1000–1099	IPX SAP access list	No
1100–1199	Extended 48-bit MAC address access list	No
1200–1299	IPX summary address access list	No

Table 19-1 Access List Numbers (continued)

Access List Number	Type	Supported
1300–1999	IP standard access list (expanded range)	Yes
2000–2699	IP extended access list (expanded range)	Yes

Note In addition to numbered standard and extended ACLs, you can also create standard and extended named IP ACLs using the supported numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Creating a Numbered Standard ACL

Beginning in privileged EXEC mode, follow these steps to create a numbered standard ACL:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*] [**log**]	Define a standard IP access list by using a source address and wildcard. The *access-list-number* is a decimal number from 1 to 99 or 1300 to 1999. Enter **deny** or **permit** to specify whether to deny or permit access if conditions are matched. The *source* is the source address of the network or host from which the packet is being sent specified as: • The 32-bit quantity in dotted-decimal format. • The keyword **any** as an abbreviation for *source* and *source-wildcard* of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard. • The keyword **host** as an abbreviation for source and source-wildcard of *source* 0.0.0.0. (Optional) The *source-wildcard* applies wildcard bits to the source. (Optional) Enter **log** to create an informational logging message about the packet that matches the entry to be sent to the console.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show access-lists** [*number* | *name*]	Show the access list configuration.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire ACL. You cannot delete individual ACEs from numbered access lists.

Note When creating an ACL, remember that, by default, the end of the ACL contains an implicit deny statement for all packets that it did not find a match for before reaching the end. With standard access lists, if you omit the mask from an associated IP host address ACL specification, 0.0.0.0 is assumed to be the mask.

This example shows how to create a standard ACL to deny access to IP host 171.69.198.102, permit access to any others, and display the results.

```
Switch (config)# access-list 2 deny host 171.69.198.102
Switch (config)# access-list 2 permit any
Switch(config)# end
Switch# show access-lists
Standard IP access list 2
    deny   171.69.198.102
    permit any
```

The switch always rewrites the order of standard access lists so that entries with **host** matches and entries with matches having a *don't care* mask of 0.0.0.0 are moved to the top of the list, above any entries with non-zero *don't care* masks. Therefore, in **show** command output and in the configuration file, the ACEs do not necessarily appear in the order in which they were entered.

The switch software can provide logging messages about packets permitted or denied by a standard IP access list. That is, any packet that matches the ACL causes an informational logging message about the packet to be sent to the console. The level of messages logged to the console is controlled by the logging console commands controlling the syslog messages.

Note Because routing is done in hardware and logging is done in software, if a large number of packets match a *permit* or *deny* ACE containing a **log** keyword, the software might not be able to match the hardware processing rate, and not all packets will be logged.

The first packet that triggers the ACL causes a logging message right away, and subsequent packets are collected over 5-minute intervals before they are displayed or logged. The logging message includes the access list number, whether the packet was permitted or denied, the source IP address of the packet, and the number of packets from that source permitted or denied in the prior 5-minute interval.

Note An output ACL cannot log multicast packets.

After creating an ACL, you must apply it to a line or interface, as described in the "Applying the ACL to an Interface or Terminal Line" section on page 19-18.

Creating a Numbered Extended ACL

Although standard ACLs use only source addresses for matching, you can use extended ACL source and destination addresses for matching operations and optional protocol type information for finer granularity of control. Some protocols also have specific parameters and keywords that apply to that protocol.

These IP protocols are supported (protocol keywords are in parentheses in bold):

Authentication Header Protocol (**ahp**), Enhanced Interior Gateway Routing Protocol (**eigrp**), Encapsulation Security Payload (**esp**), generic routing encapsulation (**gre**), Internet Control Message Protocol (**icmp**), Internet Group Management Protocol (**igmp**), Interior Gateway Routing

Protocol (**igrp**), any Interior Protocol (**ip**), IP in IP tunneling (**ipinip**), KA9Q NOS-compatible IP over IP tunneling (**nos**), Open Shortest Path First routing (**ospf**), Payload Compression Protocol (**pcp**), Protocol Independent Multicast (**pim**), Transmission Control Protocol (**tcp**), or User Datagram Protocol (**udp**).

Supported parameters can be grouped into these categories: TCP, UDP, ICMP, IGMP, or other IP.

Table 19-2 lists the possible filtering parameters for ACEs for each protocol type.

Table 19-2 Filtering Parameter ACEs Supported by Different IP Protocols

Filtering Parameter[1]	TCP	UDP	ICMP[2]	IGMP	Other IP
Layer 3 Parameters:					
IP ToS byte[3]	X	X	X	X	X
Differentiated Services Code Point (DSCP)	X	X	X	X	X
IP source address	X	X	X	X	X
IP destination address	X	X	X	X	X
Fragments	X	X	X	X	X
TCP or UDP	X	X	–	–	–
ICMP	–	–	X	–	–
IGMP	–	–	–	X	–
Other IP protocol	–	–	–	–	X
Layer 4 Parameters					
Source port	X	X	–	–	–
Source port operator	X	X	–	–	–
Destination port	X	X	–	–	–
Destination port operator	X	X	–	–	–
TCP flag	X	–	–	–	–
ICMP code	–	–	X	–	–
ICMP type	–	–	X	–	–
IGMP message type	–	–	–	X	–

1. X in a protocol column means support for the filtering parameter.
2. ICMP echo-reply cannot be filtered. All other ICMP codes or types can be filtered.
3. No support for type of service (TOS) minimize monetary cost bit.

For more details on the specific keywords relative to each protocol, refer to *Cisco IP and IP Routing Command Reference for IOS Release 12.1*.

Note The Catalyst 3550 switch does not support dynamic or reflexive access lists. It also does not support filtering based on the type of service (ToS) minimize-monetary-cost bit.

When you are creating ACEs in numbered extended access lists, remember that after you create the ACL, any additions are placed at the end of the list. You cannot reorder the list or selectively add or remove ACEs from a numbered list.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2a	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard* [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*] **Note** If you enter a **dscp** value, you cannot enter **tos** or **precedence**. You can enter both a **tos** and a **precedence** value with no **dscp**.	Define an extended IP access list and the access conditions. The *access-list-number* is a decimal number from 100 to 199 or 2000 to 2699. Enter **deny** or **permit** to specify whether to deny or permit the packet if conditions are matched. For *protocol*, enter the name or number of an IP protocol: **ahp**, **eigrp**, **esp**, **gre**, **icmp**, **igmp**, **igrp**, **ip**, **ipinip**, **nos**, **ospf**, **pcp**, **pim**, **tcp**, or **udp**, or an integer in the range 0 to 255 representing an IP protocol number. To match any Internet protocol (including ICMP, TCP, and UDP) use the keyword **ip**. **Note** This step includes options for most IP protocols. For additional specific parameters for TCP, UDP, ICMP, and IGMP, see steps 2b through 2e. The *source* is the number of the network or host from which the packet is sent. The *source-wildcard* applies wildcard bits to the source. The *destination* is the network or host number to which the packet is sent. The *destination-wildcard* applies wildcard bits to the destination. Source, source-wildcard, destination, and destination-wildcard can be specified as: • The 32-bit quantity in dotted-decimal format. • The keyword **any** for 0.0.0.0 255.255.255.255 (any host). • The keyword **host** for a single host 0.0.0.0. The other keywords are optional and have these meanings: • **precedence**—Enter to match packets with a precedence level specified as a number from 0 to 7 or by name: **routine** (0), **priority** (1), **immediate** (2), **flash** (3), **flash-override** (4), **critical** (5), **internet** (6), **network** (7). • **fragments**—Enter to check non-initial fragments. • **tos**—Enter to match by type of service level, specified by a number from 0 to 15 or a name: **normal** (0), **max-reliability** (2), **max-throughput** (4), **min-delay** (8). • **log**—Enter to create an informational logging message to be sent to the console about the packet that matches the entry or **log-input** to include the input interface in the log entry. • **time-range**—For an explanation of this keyword, see the "Applying Time Ranges to ACLs" section on page 19-15. • **dscp**—Enter to match packets with the DSCP value specified by a number from 0 to 63, or use the question mark (?) to see a list of available values.

	Command	Purpose
or	**access-list** *access-list-number* {**deny** | **permit**} *protocol* **any any** [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	In access-list configuration mode, define an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255 and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255. You can use the **any** keyword in place of source and destination address and wildcard.
or	**access-list** *access-list-number* {**deny** | **permit**} *protocol* **host** *source* **host** *destination* [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	Define an extended IP access list using an abbreviation for a source and source wildcard of *source* 0.0.0.0 and an abbreviation for a destination and destination wildcard of *destination* 0.0.0.0. You can use the **host** keyword in place of source and destination wildcard or mask.
Step 2b	**access-list** *access-list-number* {**deny** | **permit**} **tcp** *source source-wildcard* [*operator* [*port*]] *destination destination-wildcard* [*operator* [*port*]] [**established**] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*] [*flag*]	(Optional) Define an extended TCP access list and the access conditions. Enter **tcp** for Transmission Control Protocol. The parameters are the same as those described in Step 2a with these exceptions: (Optional) Enter an *operator* and *port* to compare source (if positioned after *source source-wildcard*) or destination (if positioned after *destination destination-wildcard*) port. Possible operators include **eq** (equal), **gt** (greater than), **lt** (less than), **neq** (not equal), and **range** (inclusive range). Operators require a port number (**range** requires two port numbers separated by a space). Enter the *port* number as a decimal number (from 0 to 65535) or the name of a TCP port. To see TCP port names, use the ? or refer to "Configuring IP Services" section of *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*. Use only TCP port numbers or names when filtering TCP. The additional optional keywords have these meanings: • **established**—Enter to match an established connection. This has the same function as matching on the **ack** or **rst** flag. • *flag*—Enter one of these flags to match by the specified TCP header bits: **ack** (acknowledge), **fin** (finish), **psh** (push), **rst** (reset), **syn** (synchronize), or **urg** (urgent).
Step 2c	**access-list** *access-list-number* {**deny** | **permit**} **udp** *source source-wildcard* [*operator* [*port*]] *destination destination-wildcard* [*operator* [*port*]] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Define an extended UDP access list and the access conditions. Enter **udp** for the User Datagram Protocol. The UDP parameters are the same as those described for TCP except that [*operator* [*port*]] port number or name must be a UDP port number or name, and the **flag** and **established** parameters are not valid for UDP.

	Command	Purpose
Step 2d	**access-list** *access-list-number* {**deny** | **permit**} **icmp** *source source-wildcard destination destination-wildcard* [*icmp-type* | [[*icmp-type icmp-code*] | [*icmp-message*]] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Define an extended ICMP access list and the access conditions. Enter **icmp** for Internet Control Message Protocol. The ICMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of the ICMP message type and code parameters. These optional keywords have these meanings: • *icmp-type*—Enter to filter by ICMP message type, a number from 0 to 255. • *icmp-code*—Enter to filter ICMP packets that are filtered by ICMP message type by the ICMP message code, a number from 0 to 255. • *icmp-message*—Enter to filter ICMP packets by ICMP message type name or ICMP message type and code name. To see a list of ICMP message type names and ICMP message type and code names, use the ? or refer to the "Configuring IP Services" section of *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.
Step 2e	**access-list** *access-list-number* {**deny** | **permit**} **igmp** *source source-wildcard destination destination-wildcard* [*igmp-type*] [**precedence** *precedence*] [**tos** *tos*] [**fragments**] [**log**] [**log-input**] [**time-range** *time-range-name*] [**dscp** *dscp*]	(Optional) Define an extended IGMP access list and the access conditions. Enter **igmp** for Internet Group Management Protocol. The IGMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of this optional parameter. *igmp-type*—To match IGMP message type, enter a number from 0 to 15, or enter the message name (**dvmrp**, **host-query**, **host-report**, **pim,** or **trace**).
Step 3	**show access-lists** [*number* | *name*]	Verify the access list configuration.
Step 4	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire access list. You cannot delete individual ACEs from numbered access lists.

This example shows how to create and display an extended access list to deny Telnet access from any host in network 171.69.198.0 to any host in network 172.20.52.0 and permit any others. (The **eq** keyword after the destination address means to test for the TCP destination port number equaling Telnet.)

```
Switch(config)# access-list 102 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq
telnet
Switch(config)# access-list 102 permit tcp any any
Switch(config)# end
Switch# show access-lists
Extended IP access list 102
    deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq telnet
    permit tcp any any
```

After an ACL is created, any additions (possibly entered from the terminal) are placed at the end of the list. You cannot selectively add or remove access list entries from a numbered access list.

Note When you are creating an ACL, remember that, by default, the end of the access list contains an implicit deny statement for all packets if it did not find a match before reaching the end.

After creating an ACL, you must apply it to a line or interface, as described in the "Applying the ACL to an Interface or Terminal Line" section on page 19-18.

Creating Named Standard and Extended ACLs

You can identify IP ACLs with an alphanumeric string (a name) rather than a number. You can use named ACLs to configure more IP access lists in a router than if you were to use numbered access lists. If you identify your access list with a name rather than a number, the mode and command syntax are slightly different. However, not all commands that use IP access lists accept a named access list.

Note The name you give to a standard or extended ACL can also be a number in the supported range of access list numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Consider these guidelines and limitations before configuring named ACLs:

- Not all commands that accept a numbered ACL accept a named ACL. ACLs for packet filters and route filters on interfaces can use a name. VLAN maps also accept a name.
- A standard ACL and an extended ACL cannot have the same name.
- Numbered ACLs are also available, as described in the "Creating Standard and Extended IP ACLs" section on page 19-6.
- You can apply standard and extended ACLs (named or numbered) to VLAN maps.

Beginning in privileged EXEC mode, follow these steps to create a standard ACL using names:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip access-list standard** *name*	Define a standard IP access list using a name, and enter access-list configuration mode. **Note** The name can be a number from 1 to 99.
Step 3	**deny** {*source* [*source-wildcard*] | **host** *source* | **any**} [**log**] or **permit** {*source* [*source-wildcard*] | **host** *source* | **any**} [**log**]	In access-list configuration mode, specify one or more conditions denied or permitted to determine if the packet is forwarded or dropped. • **host** *source*—A source and source wildcard of *source* 0.0.0.0. • **any**—A source and source wildcard of 0.0.0.0 255.255.255.255.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists** [*number* | *name*]	Show the access list configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a named standard ACL, use the **no ip access-list standard** *name* global configuration command.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL using names:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip access-list extended** *name*	Define an extended IP access list using a name and enter access-list configuration mode. **Note** The name can be a number from 100 to 199.
Step 3	{**deny** | **permit**} *protocol* {*source* [*source-wildcard*] | **host** *source* | **any**} {*destination* [*destination-wildcard*] | **host** *destination* | **any**} [**precedence** *precedence*] [**tos** *tos*] [**established**] [**log**] [**time-range** *time-range-name*]	In access-list configuration mode, specify the conditions allowed or denied. Use the **log** keyword to get access list logging messages, including violations. See the "Creating a Numbered Extended ACL" section on page 19-9 for definitions of protocols and other keywords. • **host** *source*—A source and source wildcard of *source* 0.0.0.0. • **host** *destination*—A destination and destination wildcard of *destination* 0.0.0.0. • **any**—A source and source wildcard or destination and destination wildcard of 0.0.0.0 255.255.255.255.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists** [*number* | *name*]	Show the access list configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a named extended ACL, use the **no ip access-list extended** *name* global configuration command.

When you are creating standard extended ACLs, remember that, by default, the end of the ACL contains an implicit deny statement for everything if it did not find a match before reaching the end. For standard ACLs, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

After you create an ACL, any additions are placed at the end of the list. You cannot selectively add ACL entries to a specific ACL. However, you can use **no permit** and **no deny** access-list configuration mode commands to remove entries from a named ACL. This example shows how you can delete individual ACEs from the named access list *border-list*:

```
Switch(config)# ip access-list extended border-list
Switch(config-ext-nacl)# no permit ip host 10.1.1.3 any
```

Being able to selectively remove lines from a named ACL is one reason you might use named ACLs instead of numbered ACLs.

After creating an ACL, you must apply it to a line or interface, as described in the "Applying the ACL to an Interface or Terminal Line" section on page 19-18.

Applying Time Ranges to ACLs

You can implement extended ACLs based on the time of day and week by using the **time-range** global configuration command. First, define the name and times of the day and week of the time range, and then reference the time range by name in an ACL to apply restrictions to the access list. You can use the time range to define when the permit or deny statements in the ACL are in effect. The **time-range** keyword

and argument are referenced in the named and numbered extended ACL task tables in the previous sections, the "Creating Standard and Extended IP ACLs" section on page 19-6, and the "Creating Named Standard and Extended ACLs" section on page 19-14.

These are some of the many possible benefits of using time ranges:

- You have more control over permitting or denying a user access to resources, such as an application (identified by an IP address/mask pair and a port number).
- You can control logging messages. ACL entries can log traffic at certain times of the day, but not constantly. Therefore, you can simply deny access without needing to analyze many logs generated during peak hours.

Note The time range relies on the switch system clock. For this feature to work the way you intend, you need a reliable clock source. We recommend that you use Network Time Protocol (NTP) to synchronize the switch clock. For more information, see the "Managing the System Time and Date" section on page 6-32.

Beginning in privileged EXEC mode, follow these steps to configure a time-range parameter for an ACL:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**time-range** *time-range-name*	Identify the time-range by a meaningful name, and enter time-range configuration mode. The name cannot contain a space or quotation mark and must begin with a letter.
Step 3	**absolute** [**start** *time date*] [**end** *time date*] or **periodic** *day-of-the-week hh:mm to* [*day-of-the-week*] *hh:mm* or **periodic** {**weekdays** | **weekend** | **daily**} *hh:mm to hh:mm*	Specify when the function it will be applied to is in effect. Use some combination of these commands; multiple **periodic** statements are allowed; only one **absolute** statement is allowed. If more than one absolute statement is configured, only the one configured last is executed.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show time-range**	Verify the time-range configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a configured time-range limitation, use the **no time-range** *time-range-name* global configuration command.

Repeat the steps if you have multiple items that you want in effect at different times. This example shows how to configure time ranges for *workhours* and for company holidays and how to verify your configuration.

```
Switch(config)# time-range workhours
Switch(config-time-range)# periodic weekdays 8:00 to 12:00
Switch(config-time-range)# periodic weekdays 13:00 to 17:00
Switch(config-time-range)# exit
Switch(config)# time-range new_year_day_2000
Switch(config-time-range)# absolute start 00:00 1 Jan 2000 end 23:59 1 Jan 2000
```

```
Switch(config-time-range)# exit
Switch(config)# time-range thanksgiving_2000
Switch(config-time-range)# absolute start 00:00 22 Nov 2000 end 23:59 23 Nov 2000
Switch(config-time-range)# exit
Switch(config)# time-range christmas_2000
Switch(config-time-range)# absolute start 00:00 24 Dec 2000 end 23:50 25 Dec 2000
Switch(config-time-range)# end
Switch# show time-range
time-range entry: christmas_2000 (inactive)
   absolute start 00:00 24 December 2000 end 23:50 25 December 2000
time-range entry: new_year_day_2000 (inactive)
   absolute start 00:00 01 January 2000 end 23:59 01 January 2000
time-range entry: thanksgiving_2000 (inactive)
   absolute start 00:00 22 November 2000 end 23:59 23 November 2000
time-range entry: workhours (inactive)
   periodic weekdays 8:00 to 12:00
   periodic weekdays 13:00 to 17:00
```

For a time range to be applied, you must reference it by name in an extended ACL that can implement time ranges. This example shows how to create and verify extended access list 188 that denies TCP traffic from any source to any destination during the defined holiday time ranges and permits all TCP traffic during work hours.

```
Switch(config)# access-list 188 deny tcp any any time-range new_year_day_2000
Switch(config)# access-list 188 deny tcp any any time-range thanskgiving_2000
Switch(config)# access-list 188 deny tcp any any time-range christmas_2000
Switch(config)# access-list 188 permit tcp any any time-range workhours
Switch(config)# end
Switch# show access-lists
Extended IP access list 188
    deny tcp any any time-range new_year_day_2000 (inactive)
    deny tcp any any time-range thanskgiving_2000 (active)
    deny tcp any any time-range christmas_2000 (inactive)
    permit tcp any any time-range workhours (inactive)
```

This example uses named ACLs to permit and deny the same traffic.

```
Switch(config)# ip access-list extended deny_access
Switch(config-ext-nacl)# deny tcp any any time-range new_year_day_2000
Switch(config-ext-nacl)# deny tcp any any time-range thanksgiving_2000
Switch(config-ext-nacl)# deny tcp any any time-range christmas_2000
Switch(config-ext-nacl)# exit
Switch(config)# ip access-list extended may_access
Switch(config-ext-nacl)# permit tcp any any time-range workhours
Switch(config-ext-nacl)# end
Switch# show ip access-lists
Extended IP access list deny_access
    deny tcp any any time-range new_year_day_2000 (inactive)
    deny tcp any any time-range thanksgiving_2000 (inactive)
    deny tcp any any time-range christmas_2000 (inactive)
Extended IP access list may_access
    permit tcp any any time-range workhours (inactive)
```

Including Comments About Entries in ACLs

You can use the **remark** keyword to include comments (remarks) about entries in any IP standard or extended ACL. The remarks make the ACL easier for you to understand and scan. Each remark line is limited to 100 characters.

The remark can go before or after a permit or deny statement. You should be consistent about where you put the remark so that it is clear which remark describes which permit or deny statement. For example, it would be confusing to have some remarks before the associated permit or deny statements and some remarks after the associated statements.

For IP numbered standard or extended ACLs, use the **access-list** *access-list number* **remark** *remark* global configuration command to include a comment about an access list. To remove the remark, use the **no** form of this command.

In this example, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

For an entry in a named IP ACL, use the **remark** access-list configuration command. To remove the remark, use the **no** form of this command.

In this example, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp host 171.69.2.88 any eq telnet
```

Applying the ACL to an Interface or Terminal Line

After you create an ACL, you can apply it to one or more interfaces or terminal lines. ACLs can be applied on *either* outbound or inbound interfaces. This section describes how to accomplish this task for both terminal lines and network interfaces. Note these guidelines:

- When controlling access to a line, you must use a number. Only numbered ACLs can be applied to lines.
- When controlling access to an interface, you can use a name or number.
- Set identical restrictions on all the virtual terminal lines because a user can attempt to connect to any of them.
- If you apply an ACL to a Layer-3 interface and the enhanced multilayer software image is not installed on your switch, the ACL only filters packets that are intended for the CPU, such as SNMP, Telnet, or Web traffic.

Beginning in privileged EXEC mode, follow these steps to restrict incoming and outgoing connections between a virtual terminal line and the addresses in an ACL:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**line** [**console** | **vty**] *line-number*	Identify a specific line for configuration, and enter in-line configuration mode. • **console**—Enter to specify the console terminal line. The console port is DCE. • **vty**—Enter to specify a virtual terminal for remote console access. The *line-number* is the first line number in a contiguous group that you want to configure when the line type is specified. The range is from 0 to 16.
Step 3	**access-class** *access-list-number* {**in** | **out**}	Restrict incoming and outgoing connections between a particular virtual terminal line (into a device) and the addresses in an access list.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Display the access list configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove access restrictions on a terminal line, use the **no access-class** *access-list-number* {**in** | **out**} line configuration command.

Beginning in privileged EXEC mode, follow these steps to control access to a Layer 3 interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Identify a specific interface for configuration, and enter interface configuration mode. The interface must be a Layer 3 interface, either a routed port or an SVI VLAN ID.
Step 3	**ip access-group** {*access-list-number* | *name*} {**in** | **out**}	Control access to the specified interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Display the access list configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the specified access group, use the **no ip access-group** {*access-list-number* | *name*} {**in** | **out**} interface configuration command.

This example shows how to apply access list 2 on Gigabit Ethernet interface 0/3 to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet0/3
Router(config-if)# ip access-group 2 in
```

Note The **ip access-group** interface configuration command is only valid when applied to a Layer 3 interface: an SVI, a Layer 3 EtherChannel, or a routed port. The interface must have been configured with an IP address. Layer 3 access groups filter packets that are routed or are received by Layer 3 processes on the CPU. They do not affect packets bridged within a VLAN.

For inbound ACLs, after receiving a packet, the switch checks the packet against the ACL. If the ACL permits the packet, the switch continues to process the packet. If the ACL rejects the packet, the switch discards the packet.

For outbound ACLs, after receiving and routing a packet to a controlled interface, the switch checks the packet against the ACL. If the ACL permits the packet, the switch sends the packet. If the ACL rejects the packet, the switch discards the packet.

If the input interface is configured to send ICMP Unreachable messages, these messages are sent whenever a packet is discarded, regardless of whether the packet was discarded because of an ACL on the input interface or because of an ACL on the output interface. ICMP Unreachables are normally limited to no more than one every one-half second per input interface, but this can be changed by using the **ip icmp rate-limit unreachable** global configuration command.

When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied to the interface and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Displaying ACLs and Access Groups

You can display existing ACLs and when you use the **ip access-group** interface configuration command to apply ACLs to a Layer 3 interface, you can display the access groups on the interface. You can use the privileged EXEC commands as described in Table 19-3 to display this information.

Table 19-3 Commands for Displaying Access Lists and Access Groups

Command	Purpose
show access-lists [*number* | *name*]	Display the contents of one or all current IP and MAC address access lists or a specific access list (numbered or named).
show ip access-lists [*number* | *name*]	Display the contents of all current IP access lists or a specific IP access list (numbered or named).
show ip interface *interface-id*	Display detailed configuration and status of an interface. If IP is enabled on the interface and ACLs have been applied by using the **ip access-group** interface configuration command, the access groups are included in the display.
show running-config [**interface** *interface-id*]	Displays the contents of the configuration file for the switch or the specified interface, including all configured access groups, whether or not IP is enabled on an interface.

This is an example of output from the **show access-lists** privileged EXEC command, displaying all standard and extended ACLs:

```
Switch# show access-lists
Standard IP access list 1
    permit 172.20.10.10
Standard IP access list 10
    permit 12.12.12.12
Standard IP access list 12
    deny   1.3.3.2
Standard IP access list 32
    permit 172.20.20.20
Standard IP access list 34
    permit 10.24.35.56
    permit 23.45.56.34
Extended IP access list 120
    permit eigrp host 12.3.6.5 host 25.36.1.24
Extended MAC access list mac1
```

This is an example of output from the **show ip access-lists** privileged EXEC command. It displays only IP standard and extended ACLs. Note that the named MAC extended ACL displayed in the previous example is not included in this display.

```
Switch# show ip access-lists
Standard IP access list 1
    permit 172.20.10.10
Standard IP access list 10
    permit 12.12.12.12
Standard IP access list 12
    deny   1.3.3.2
Standard IP access list 32
    permit 172.20.20.20
Standard IP access list 34
    permit 10.24.35.56
    permit 23.45.56.34
Extended IP access list 120
    permit eigrp host 12.3.6.5 host 25.36.1.24
```

This example shows how to view all access groups configured for Gigabit Ethernet interface 0/2, which has IP enabled:

```
Switch# show ip interface gigabitethernet0/2
GigabitEthernet0/2 is up, line protocol is down
  Internet address is 10.20.30.1/16
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is permit Any
  Inbound  access list is 13

<output truncated>
```

This example shows how to use the **show running-config interface** privileged EXEC command to display the ACL configuration of Gigabit Ethernet interface 0/2:

```
Switch# show running-config interface gigabitethernet0/2
Building configuration...

Building configuration...

Current configuration : 85 bytes
!
interface GigabitEthernet0/2
 ip address 10.20.30.1 255.255.0.0
 ip access-group 13 in
 ip access-group permit Any out
 no switchport
!
<output truncated>
!
access-list 13 permit any log
access-list 101 permit icmp any any conversion-error
access-list 101 permit 234 host 172.30.40.1 host 123.23.23.2
access-list 103 permit icmp any any 123 23 tos max-throughput
access-list 103 permit igmp any any 12
<information truncated>
!
```

ACL Configuration Examples

This section provides examples of configuring ACLs. For detailed information about compiling ACLs, refer to the *Security Configuration Guide* and the "IP Services" chapter of the *Cisco IOS IP and IP Routing Configuration Guide for IOS Release 12.1*.

Figure 19-3 shows a small networked office environment with the routed port 0/2 connected to Server A, containing benefits and other information that all employees can access, and routed port 0/3 connected to Server B, containing confidential payroll data. All users can access Server A, but Server B has restricted access.

Use router ACLs to do this in one of two ways:

- Create a standard ACL, and filter traffic coming to the server from port 0/3.
- Create an extended ACL, and filter traffic coming from the server into port 0/3.

Figure 19-3 Using Router ACLs to Control Traffic

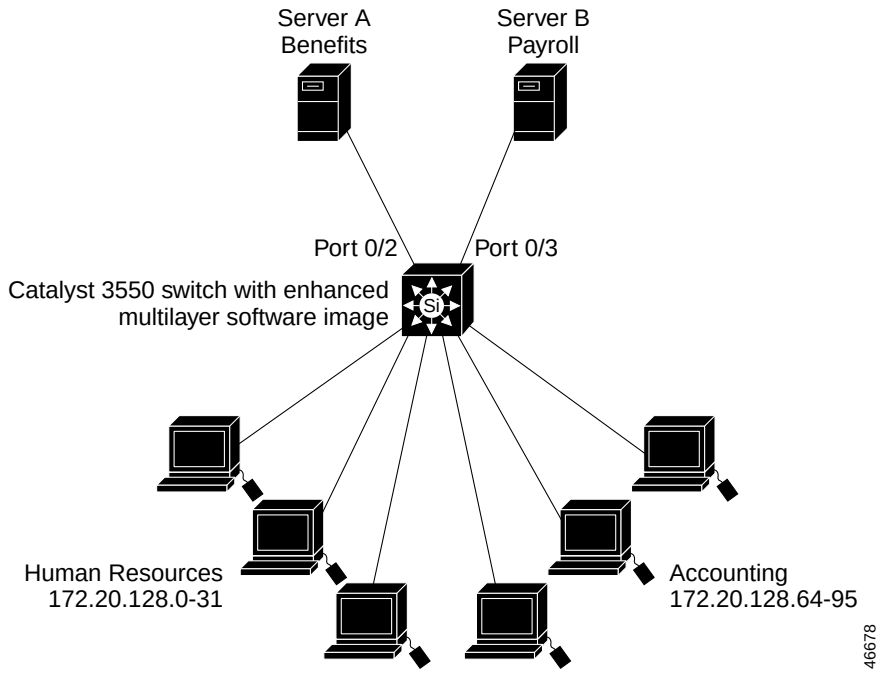

This example uses a standard ACL to filter traffic coming into Server B from port 0/3, permitting traffic only from Accounting's source addresses 172.20.128.64 to 172.20.128.95.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0.31
Switch(config)# end
Switch# show access-lists
Standard IP access list 6
    permit 172.20.128.64, wildcard bits 0.0.0.31
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# ip access-group 6 out
```

The ACL is applied to traffic coming out of routed port 0/3 from the specified source address.

This example uses an extended ACL to filter traffic coming from Server B into port 0/3, permitting traffic from any source address (in this case Server B) to only the Accounting destination addresses 172.20.128.64 to 172.20.128.95.

```
Switch(config)# access-list 106 permit ip any 172.20.128.64 0.0.0.31
Switch(config)# end
Switch# show access-lists
Extended IP access list 106
    permit ip any 172.20.128.64 0.0.0.31
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# ip access-group 106 in
```

The ACL is then applied to traffic going into routed port 0/3, permitting it to go only to the specified destination addresses. Note that with extended ACLs, you must enter the protocol (IP) before the source and destination information.

Numbered ACLs

In this example, network 36.0.0.0 is a Class A network whose second octet specifies a subnet; that is, its subnet mask is 255.255.0.0. The third and fourth octets of a network 36.0.0.0 address specify a particular host. Using access list 2, the switch accepts one address on subnet 48 and reject all others on that subnet. The last line of the list shows that the switch accepts addresses on all other network 36.0.0.0 subnets. The ACL is then applied to packets entering Gigabit Ethernet interface 0/1.

```
Switch(config)# access-list 2 permit 36.48.0.3
Switch(config)# access-list 2 deny 36.48.0.0 0.0.255.255
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 2 in
```

Extended ACLs

In this example, the first line permits any incoming TCP connections with destination ports greater than 1023. The second line permits incoming TCP connections to the Simple Mail Transfer Protocol (SMTP) port of host 128.88.1.2. The third line permits incoming ICMP messages for error feedback.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 gt 1023
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# access-list 102 permit icmp any any
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 102 in
```

For another example of using an extended ACL, suppose that you have a network connected to the Internet, and you want any host on the network to be able to form TCP connections to any host on the Internet. However, you do not want IP hosts to be able to form TCP connections to hosts on your network, except to the mail (SMTP) port of a dedicated mail host.

SMTP uses TCP port 25 on one end of the connection and a random port number on the other end. The same port numbers are used throughout the life of the connection. Mail packets coming in from the Internet have a destination port of 25. Outbound packets have the port numbers reversed. Because the secure system behind the router always accepts mail connections on port 25, the incoming and outgoing services are separately controlled. The ACL must be configured as an input ACL on the outbound interface and an output ACL on the inbound interface.

In this example, the network is a Class B network with the address 128.88.0.0, and the mail host address is 128.88.1.2. The **established** keyword is used only for the TCP to show an established connection. A match occurs if the TCP datagram has the ACK or RST bits set, which show that the packet belongs to an existing connection. Gigabit Ethernet interface 0/1 is the interface that connects the router to the Internet.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 established
Switch(config)# access-list 102 permit tcp any host 128.88.1.2 eq 25
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 102 in
```

Named ACLs

The following configuration creates a standard ACL named *internet_filter* and an extended ACL named *marketing_group*. The *internet_filter* ACL allows all traffic from the source address 1.2.3.4.

```
Switch(config)# ip access-list standard Internet_filter
Switch(config-ext-nacl)# permit 1.2.3.4
Switch(config-ext-nacl)# exit
```

The *marketing_group* ACL allows any TCP Telnet traffic to the destination address and wildcard 171.69.0.0 0.0.255.255 and denies any other TCP traffic. It permits any ICMP traffic, denies UDP traffic from any source to the destination address range 171.69.0.0 through 179.69.255.255 with a destination port less than 1024, denies any other IP traffic, and provides a log of the result.

```
Switch(config)# ip access-list extended marketing_group
Switch(config-ext-nacl)# permit tcp any 171.69.0.0 0.0.255.255 eq telnet
Switch(config-ext-nacl)# deny tcp any any
Switch(config-ext-nacl)# permit icmp any any
Switch(config-ext-nacl)# deny udp any 171.69.0.0 0.0.255.255 lt 1024
Switch(config-ext-nacl)# deny ip any any log
Switch(config-ext-nacl)# exit
```

The ACLs are applied to Gigabit Ethernet port 0/5, which is configured as a Layer 3 port, with the *Internet_filter* ACL applied to incoming traffic and the *marketing_group* ACL applied to outgoing traffic.

```
Switch(config)# interface gigabitethernet0/5
Switch(config-if)# no switchport
Switch(config-if)# ip address 2.0.5.1 255.255.255.0
Switch(config-if)# ip access-group Internet_filter out
Switch(config-if)# ip access-group marketing_group in
...
```

Time Range Applied to an IP ACL

This example denies Hypertext Transfer Protocol (HTTP) traffic on IP on Monday through Friday between the hours of 8:00 a.m. and 6:00 p.m. The example allows UDP traffic only on Saturday and Sunday from noon to 8:00 p.m.

```
Switch(config)# time-range no-http
Switch(config)# periodic weekdays 8:00 to 18:00
!
Switch(config)# time-range udp-yes
Switch(config)# periodic weekend 12:00 to 20:00
!
Switch(config)# ip access-list extended strict
Switch(config-ext-nacl)# deny tcp any any eq www time-range no-http
Switch(config-ext-nacl)# permit udp any any time-range udp-yes
!
Switch(config-ext-nacl)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group strict in
```

Commented IP ACL Entries

In this example of a numbered ACL, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

In this example of a numbered ACL, the Winter and Smith workstations are not allowed to browse the Web:

```
Switch(config)# access-list 100 remark Do not allow Winter to browse the web
Switch(config)# access-list 100 deny host 171.69.3.85 any eq www
Switch(config)# access-list 100 remark Do not allow Smith to browse the web
Switch(config)# access-list 100 deny host 171.69.3.13 any eq www
```

In this example of a named ACL, the Jones subnet is not allowed access:

```
Switch(config)# ip access-list standard prevention
Switch(config-std-nacl)# remark Do not allow Jones subnet through
Switch(config-std-nacl)# deny 171.69.0.0 0.0.255.255
```

In this example of a named ACL, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp 171.69.0.0 0.0.255.255 any eq telnet
```

ACL Logging

Two variations of logging are supported on router ACLs. The **log** keyword sends an informational logging message to the console about the packet that matches the entry; the **log-input** keyword includes the input interface in the log entry.

In this example, standard named access list *stan1* denies traffic from 10.1.1.0 0.0.0.255, allows traffic from all other sources, and includes the **log** keyword.

```
Switch(config)# ip access-list standard stan1
Switch(config-std-nacl)# deny 10.1.1.0 0.0.0.255 log
Switch(config-std-nacl)# permit any log
Switch(config-std-nacl)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group stan1 in
Switch(config-if)# end

Switch# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
    Console logging: level debugging, 37 messages logged
    Monitor logging: level debugging, 0 messages logged
    Buffer logging: level debugging, 37 messages logged
    File logging: disabled
    Trap logging: level debugging, 39 message lines logged

Log Buffer (4096 bytes):

00:00:48: NTP: authentication delay calculation problems

<output truncated>

00:09:34:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
00:09:59:%SEC-6-IPACCESSLOGS:list stan1 denied 10.1.1.15 1 packet
00:10:11:%SEC-6-IPACCESSLOGS:list stan1 permitted 0.0.0.0 1 packet
00:15:33:%SEC-6-IPACCESSLOGS:list stan1 denied 10.1.1.15 2009 packets
```

This example is a named extended access list *ext1* that permits ICMP packets from any source to 10.1.1.0 0.0.0.255 and denies all UDP packets.

```
Switch(config)# ip access-list extended ext1
Switch(config-ext-nacl)# permit icmp any 10.1.1.0 0.0.0.255 log
Switch(config-ext-nacl)# deny udp any any log
Switch(config-std-nacl)# exit
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# ip access-group ext1 in
```

This is a an example of a log for an extended ACL:

```
01:24:23:%SEC-6-IPACCESSLOGDP:list ext1 permitted icmp 10.1.1.15 -> 10.1.1.61 (0/0), 1
packet
01:25:14:%SEC-6-IPACCESSLOGDP:list ext1 permitted icmp 10.1.1.15 -> 10.1.1.61 (0/0), 7
packets
01:26:12:%SEC-6-IPACCESSLOGP:list ext1 denied udp 0.0.0.0(0) -> 255.255.255.255(0), 1
packet
01:31:33:%SEC-6-IPACCESSLOGP:list ext1 denied udp 0.0.0.0(0) -> 255.255.255.255(0), 8
packets
```

Note that all logging entries for IP ACLs start with `%SEC-6-IPACCESSLOG` with minor variations in format depending on the kind of ACL and the access entry that has been matched.

This is an example of an output message when the **log-input** keyword is entered:

```
00:04:21:%SEC-6-IPACCESSLOGDP:list inputlog permitted icmp 10.1.1.10 (Vlan1
0001.42ef.a400) -> 10.1.1.61 (0/0), 1 packet
```

A log message for the same sort of packet using the **log** keyword does not include the input interface information:

```
00:05:47:%SEC-6-IPACCESSLOGDP:list inputlog permitted icmp 10.1.1.10 -> 10.1.1.61 (0/0), 1
packet
```

Configuring VLAN Maps

This section describes how to configure VLAN maps, which is the only way to control filtering within a VLAN. VLAN maps have no direction. To filter traffic in a specific direction by using a VLAN map, you need to include an ACL with specific source or destination addresses. If there is a match clause for that type of packet (IP or MAC) in the VLAN map, the default action is to drop the packet if the packet does not match any of the entries within the map. If there is no match clause for that type of packet, the default is to forward the packet.

Note For complete syntax and usage information for the commands used in this section, refer to the Catalyst *3550 Multilayer Switch Command Reference* for this release.

To create a VLAN map and apply it to one or more VLANs, perform these steps:

Step 1 Create the standard or extended IP ACLs or named MAC extended ACLs that you want to apply to the VLAN. See the "Creating Standard and Extended IP ACLs" section on page 19-6 and the "Creating Named MAC Extended ACLs" section on page 19-28.

Step 2 Enter the **vlan access-map** global configuration command to create a VLAN ACL map entry.

Step 3 In access map configuration mode, optionally enter an **action**—**forward** (the default) or **drop**—and enter the **match** command to specify an IP packet or a non-IP packet (with only a known MAC address) and to match the packet against one or more ACLs (standard or extended).

Note If the VLAN map has a match clause for the type of packet (IP or MAC) and the packet does not match the type, the default is to drop the packet. If there is no match clause in the VLAN map for that type of packet, and no action specified, the packet is forwarded.

Step 4 Use the **vlan filter** global configuration command to apply a VLAN map to one or more VLANs.

This section contains these topics:

- VLAN Map Configuration Guidelines, page 19-28
- Creating Named MAC Extended ACLs, page 19-28
- Creating a VLAN Map, page 19-30
- Applying a VLAN Map to a VLAN, page 19-32
- Displaying VLAN Map Information, page 19-33

VLAN Map Configuration Guidelines

Follow these guidelines when configuring VLAN maps:

- If there is no router ACL configured to deny traffic on a routed VLAN interface (input or output), and *no* VLAN map configured, all traffic is permitted.
- Each VLAN map consists of a series of entries. The order of entries in an VLAN map is important. A packet that comes into the switch is tested against the first entry in the VLAN map. If it matches, the action specified for that part of the VLAN map is taken. If there is no match, the packet is tested against the next entry in the map.
- If the VLAN map has at least one match clause for the type of packet (IP or MAC) and the packet does not match any of these match clauses, the default is to drop the packet. If there is no match clause for that type of packet in the VLAN map, the default is to forward the packet.
- The system might take longer to boot if you have configured a very large number of ACLs.
- For information about using both router ACLs and VLAN maps, see the "Guidelines" section on page 19-36.
- See the "Using VLAN Maps in Your Network" section on page 19-33 for configuration examples.

Creating Named MAC Extended ACLs

You can filter non-IP traffic on a VLAN by using MAC addresses and named MAC extended ACLs. The procedure is similar to that of configuring other extended named ACLs.

Note Named MAC extended ACLs can only be applied to VLAN maps.

For more information about the supported non-IP protocols in the **mac access-list extended** command, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Note Though visible in the command-line help strings, **appletalk** is not supported as a matching condition for the **deny** and **permit** MAC access-list configuration mode commands, nor is matching on any SNAP-encapsulated packet with a non-zero Organizational Unique Identifier (OUI).

Beginning in privileged EXEC mode, follow these steps to create a named MAC extended ACL:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mac access-list extended** *name*	Define an extended MAC access list using a name.
Step 3	{**deny** | **permit**} {**any** | **host** *source MAC address* | *source MAC address mask*} {**any** | **host** *destination MAC address* | *destination MAC address mask*} [*type mask* | **lsap** *lsap mask* | **aarp** | **amber** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp** | *0-65535*] [**cos** *cos*]	In extended MAC access-list configuration mode, specify to **permit** or **deny any** source MAC address, a source MAC address with a mask, or a specific **host** source MAC address and **any** destination MAC address, destination MAC address with a mask, or a specific destination MAC address. (Optional) You can also enter these options: • *type mask*—An arbitrary EtherType number of a packet with Ethernet II or SNAP encapsulation in decimal, hex, or octal with optional mask of *don't care* bits applied to the EtherType before testing for a match. • **lsap** *lsap mask*—An LSAP number of a packet with 802.2 encapsulation in decimal, hex, or octal with optional mask of *don't care* bits. • **aarp** | **amber** | **dec-spanning** | **decnet-iv** | **diagnostic** | **dsm** | **etype-6000** | **etype-8042** | **lat** | **lavc-sca** | **mop-console** | **mop-dump** | **msdos** | **mumps** | **netbios** | **vines-echo** |**vines-ip** | **xns-idp**—A non-IP protocol. • **cos** *cos*—An IEEE 802.1Q cost of service number from 0 to 7 used to set priority.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists** [*number* | *name*]	Show the access list configuration.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no mac access-list extended** *name* global configuration command to delete the entire ACL. You can also delete individual ACEs from named MAC extended ACLs.

This example shows how to create and display an access list named *mac1*, denying only EtherType DECnet Phase IV traffic, but permitting all other types of traffic.

```
Switch(config)# mac access-list extended mac1
Switch(config-ext-macl)# deny any any decnet-iv
Switch(config-ext-macl)# permit any any
Switch(config-ext-macl)# end
Switch # show access-lists
Extended MAC access list mac1
    deny   any any decnet-iv
    permit any any
```

Creating a VLAN Map

Each VLAN map consists of an ordered series of entries. Beginning in privileged EXEC mode, follow these steps to create, add to, or delete a VLAN map entry:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**vlan access-map** *name* [*number*]	Create a VLAN map, and give it a name and (optionally) a number. The number is the sequence number of the entry within the map. When you create VLAN maps with the same name, numbers are assigned sequentially in increments of 10. When modifying or deleting maps, you can enter the number of the map entry that you want to modify or delete. Entering this command changes to access-map configuration mode.
Step 3	**action** {**drop** | **forward**}	(Optional) Set the action for the map entry. The default is to forward.
Step 4	**match** {**ip** | **mac**} **address** {*name* | *number*} [*name* | *number*]	Match the packet (using either the IP or MAC address) against one or more standard or extended access lists. Note that packets are only matched against access lists of the correct protocol type. IP packets are matched against standard or extended IP access lists. Non-IP packets are only matched against named MAC extended access lists.
Step 5	**end**	Return to global configuration mode.
Step 6	**show running-config**	Display the access list configuration.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no vlan access-map** *name* global configuration command to delete a map.

Use the **no vlan access-map** *name number* global configuration command to delete a single sequence entry from within the map.

Use the **no action** access-map configuration command to enforce the default action, which is to forward.

VLAN maps do not use the specific permit or deny keywords. To deny a packet by using VLAN maps, create an ACL that would match the packet, and set the action to drop. A permit in the ACL counts as a match. A deny in the ACL means no match.

Examples of ACLs and VLAN Maps

These examples show how to create ACLs and VLAN maps that for specific purposes.

Example 1

This example shows how to create an ACL and a VLAN map to deny a packet. In the first map, any packets that match the *ip1* ACL (TCP packets) would be dropped. You first create the *ip1*ACL to permit any TCP packet and no other packets. Because there is a match clause for IP packets in the VLAN map, the default action is to drop any IP packet that does not match any of the match clauses.

```
Switch(config)# ip access-list extended ip1
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 10
Switch(config-access-map)# match ip address ip1
Switch(config-access-map)# action drop
```

This example shows how to create a VLAN map to permit a packet. ACL *ip2* permits UDP packets and any packets that match the *ip2* ACL are forwarded.

```
Switch(config)# ip access-list extended ip2
Switch(config-ext-nacl)# permit udp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 20
Switch(config-access-map)# match ip address ip2
Switch(config-access-map)# action forward
```

In this map, any IP packets that did not match any of the previous ACLs (that is, packets that are not TCP packets or UDP packets) would get dropped.

Example 2

In this example, the VLAN map has a default action of drop for IP packets and a default action of forward for MAC packets. Used with standard ACL 101 and extended named access lists **igmp-match** and **tcp-match**, the map will have the following results:

- Forward all UDP packets
- Drop all IGMP packets
- Forward all TCP packets
- Drop all other IP packets
- Forward all non-IP packets

```
Switch(config)# access-list 101 permit udp any any
Switch(config)# ip access-list extended igmp-match
Switch(config-ext-nacl)# permit igmp any any
Switch(config)# ip access-list extended tcp-match
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map drop-ip-default 10
Switch(config-access-map)# match ip address 101
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 20
Switch(config-access-map)# match ip address igmp-match
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 30
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
```

Example 3

In this example, the VLAN map has a default action of drop for MAC packets and a default action of forward for IP packets. Used with MAC extended access lists **good-hosts** and **good-protocols**, the map will have the following results:

- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Forward MAC packets with decnet-iv or vines-ip protocols
- Drop all other non-IP packets
- Forward all IP packets

```
Switch(config)# mac access-list extended good-hosts
Switch(config-ext-macl)# permit host 000.0c00.0111 any
Switch(config-ext-macl)# permit host 000.0c00.0211 any
Switch(config-ext-nacl)# exit
Switch(config)# mac access-list extended good-protocols
Switch(config-ext-macl)# permit any any decnet-ip
Switch(config-ext-macl)# permit any any vines-ip
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map drop-mac-default 10
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-mac-default 20
Switch(config-access-map)# match mac address good-protocols
Switch(config-access-map)# action forward
```

Example 4

In this example, the VLAN map has a default action of drop for all packets (IP and non-IP). Used with access lists **tcp-match** and **good-hosts** from Examples 2 and 3, the map will have the following results:

- Forward all TCP packets
- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Drop all other IP packets
- Drop all other MAC packets

```
Switch(config)# vlan access-map drop-all-default 10
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-all-default 20
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
```

Applying a VLAN Map to a VLAN

Beginning in privileged EXEC mode, follow these steps to apply a VLAN map to one or more VLANs:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**vlan filter** *mapname* **vlan-list** *list*	Apply the VLAN map to one or more VLAN IDs. The list can be a single VLAN ID (22), a consecutive list (10-22), or a string of VLAN IDs (12, 22, 30). Spaces around the comma and hyphen are optional.
Step 3	**show running-config**	Display the access list configuration.
Step 4	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the VLAN map, use the **no vlan filter** *mapname* **vlan-list** *list* global configuration command.

This example shows how to apply VLAN map 1 to VLANs 20 through 22:

```
Switch(config)# vlan filter map 1 vlan-list 20-22
```

Displaying VLAN Map Information

You can display information about VLAN access maps or VLAN filters. Use the privileged EXEC commands in Table 19-4 to display VLAN map information.

Table 19-4 Commands for Displaying VLAN Map Information

Command	Purpose
show vlan access-map [*mapname*]	Show information about all VLAN access-maps or the specified access map.
show vlan filter [**access-map** *name* | **vlan** *vlan-id*]	Show information about all VLAN filters or about a specified VLAN or VLAN access map.

This is an example of output from the **show vlan access-map** privileged EXEC command:

```
Switch# show vlan access-map
Vlan access-map "map_1" 10
  Match clauses:
     ip address: ip1
  Action:
     drop
Vlan access-map "map_1" 20
  Match clauses:
     mac address: mac1
  Action:
     forward
Vlan access-map "map_1" 30
  Match clauses:
  Action:
     drop
```

This is an example of output from the **show vlan filter** privileged EXEC command:

```
Switch# show vlan filter
VLAN Map map_1 is filtering VLANs:
  20-22
```

Using VLAN Maps in Your Network

This section describes some typical uses for VLAN maps and includes these topics:

- Wiring Closet Configuration, page 19-34
- Denying Access to a Server on Another VLAN, page 19-35

Wiring Closet Configuration

In a wiring closet configuration, the Catalyst 3550 switch might not be running the enhanced multilayer software image. In this configuration, the switch can still support a VLAN map and a QoS classification ACL. In Figure 19-4, assume that Host X and Host Y are in different VLANs and are connected to wiring closet switches A and C. Traffic from Host X to Host Y is eventually being routed by Switch B, which is running the enhanced multilayer software image. Traffic from Host X to Host Y can be access-controlled at the traffic entry point, Switch A.

Figure 19-4 Wiring Closet Configuration

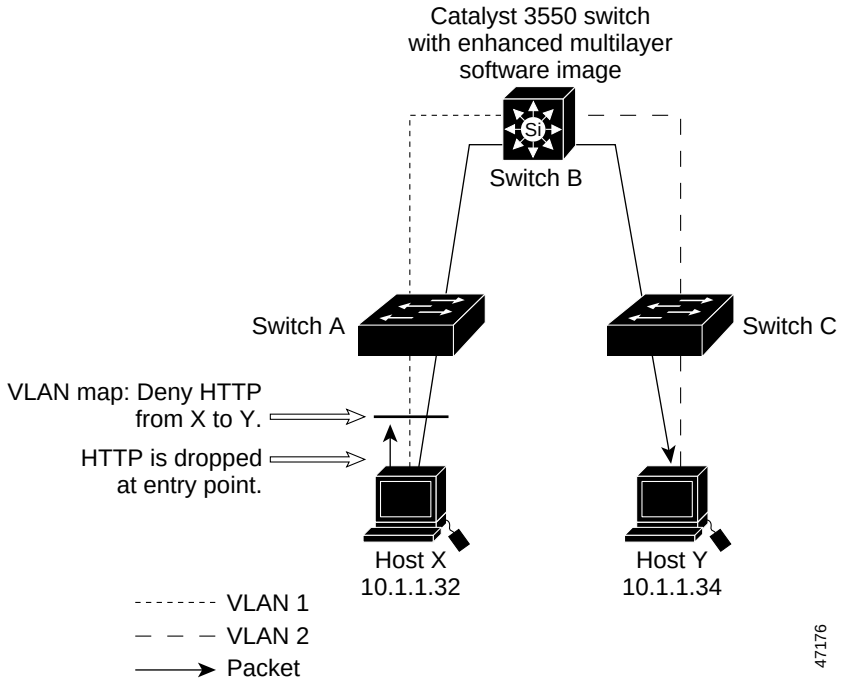

If you do not want HTTP traffic switched from Host X to Host Y, you can configure a VLAN map on Switch A to drop all HTTP traffic from Host X (IP address 10.1.1.32) to Host Y (IP address 10.1.1.34) at Switch A and not bridge it to Switch B.

First, define the IP access list *http* that permits (matches) any TCP traffic on the HTTP port.

```
Switch(config)# ip access-list extended http
Switch(config-ext-nacl)# permit tcp host 10.1.1.32 host 10.1.1.34 eq www
Switch(config-ext-nacl)# exit
```

Next, create VLAN access map *map2* so that traffic that matches the *http* access list is dropped and all other IP traffic is forwarded.

```
Switch(config)# vlan access-map map2 10
Switch(config-access-map)# match ip address http
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# ip access-list extended match_all
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map2 20
Switch(config-access-map)# match ip address match_all
Switch(config-access-map)# action forward
```

Then, apply VLAN access map *map2*to VLAN 1.

```
Switch(config)# vlan filter map2 vlan 1
```

Denying Access to a Server on Another VLAN

You can restrict access to a server on another VLAN. For example, server 10.1.1.100 in VLAN 10 needs to have access restricted as follows (see Figure 19-5):

- Hosts in subnet 10.1.2.0/8 in VLAN 20 should not have access.
- Hosts 10.1.1.4 and 10.1.1.8 in VLAN 10 should not have access.

Figure 19-5 Deny Access to a Server on Another VLAN

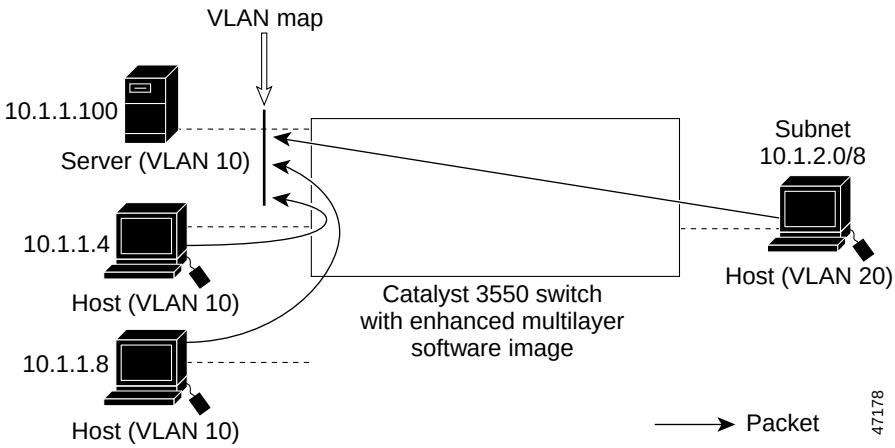

This example shows how to deny access to a server on another VLAN by creating the VLAN map SERVER 1 that denies access to hosts in subnet 10.1.2.0/8, host 10.1.1.4, and host 10.1.1.8 and permits other IP traffic. The final step is to apply the map SERVER1 to VLAN 10.

Step 1 Define the IP ACL that will match the correct packets.

```
Switch(config)# ip access-list extended SERVER1_ACL
Switch(config-ext-nacl))# permit ip 10.1.2.0 0.0.0.255 host 10.1.1.100
Switch(config-ext-nacl))# permit ip host 10.1.1.4 host 10.1.1.100
Switch(config-ext-nacl))# permit ip host 10.1.1.8 host 10.1.1.100
Switch(config-ext-nacl))# exit
```

Step 2 Define a VLAN map using this ACL that will drop IP packets that match SERVER1_ACL and forward IP packets that do not match the ACL.

```
Switch(config)# vlan access-map SERVER1_MAP
Switch(config-access-map)# match ip address SERVER1_ACL
Switch(config-access-map)# action drop
Switch(config)# vlan access-map SERVER1_MAP 20
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
```

Step 3 Apply the VLAN map to VLAN 10.

```
Switch(config)# vlan filter SERVER1_MAP vlan-list 10.
```

Using VLAN Maps with Router ACLs

To access control both bridged and routed traffic, you can use VLAN maps only or a combination of router ACLs and VLAN maps. You can define router ACLs on both input and output routed VLAN interfaces, and you can define a VLAN map to access control the bridged traffic.

If a packet flow matches a VLAN-map deny clause in the ACL, regardless of the router ACL configuration, the packet flow is denied.

Note When you use router ACLs with VLAN maps, packets that require logging on the router ACLs are not logged if they are denied by a VLAN map.

If the VLAN map has a match clause for the type of packet (IP or MAC) and the packet does not match the type, the default is to drop the packet. If there is no match clause in the VLAN map, and no action specified, the packet is forwarded if it does not match any VLAN map entry.

This section includes this information about using VLAN maps with router ACLs:

- Guidelines, page 19-36
- Determining if the ACL Configuration Fits in Hardware, page 19-37
- Examples of Router ACLs and VLAN Maps Applied to VLANs, page 19-39

Guidelines

These guidelines are for configurations where you need to have an router ACL *and* a VLAN map on the same VLAN. These guidelines do not apply to configurations where you are mapping router ACLs and VLAN maps on different VLANs.

The switch hardware provides one lookup for security ACLs for each direction (input and output); therefore, you must merge a router ACL and a VLAN map when they are configured on the same VLAN. Merging the router ACL with the VLAN map might significantly increase the number of ACEs.

If you must configure a router ACL and a VLAN map on the same VLAN, use these guidelines for both router ACL and VLAN map configuration:

- Whenever possible, try to write the ACL with all entries having a single action except for the final, default action of the other type. That is, write the ACL using one of these two forms:

  permit...
  permit...
  permit...
  deny ip any any

  or

  deny...
  deny...
  deny...
  permit ip any any

- To define multiple actions in an ACL (permit, deny), group each action type together to reduce the number of entries.

- Avoid including Layer 4 information in an ACL; adding this information complicates the merging process. The best merge results are obtained if the ACLs are filtered based on IP addresses (source and destination) and not on the full flow (source IP address, destination IP address, protocol, and protocol ports). It is also helpful to use *don't care* bits in the IP address, whenever possible.

  If you need to specify the full-flow mode and the ACL contains both IP ACEs and TCP/UDP/ICMP ACEs with Layer 4 information, put the Layer 4 ACEs at the end of the list. This gives priority to the filtering of traffic based on IP addresses.

Determining if the ACL Configuration Fits in Hardware

As previously stated, ACL processing in the Catalyst 3550 switch is mostly accomplished in hardware. However, if the hardware reaches its capacity to store ACL configurations, the switch software attempts to fit a simpler configuration into the hardware. This simpler configuration does not do all the filtering that has been configured, but instead sends some or all packets to the CPU to be filtered by software. In this way, all configured filtering will be accomplished, but performance is greatly decreased when the filtering is done in software.

For example, if the combination of an input router ACL applied to a VLAN interface and a VLAN map applied to the same VLAN does not fit into the hardware, these results might occur:

- If the VLAN map alone fits in hardware, the software sets up the hardware to send to the CPU all packets that need to be routed for filtering and possible routing (if the packet passes the filter). Packets that only require bridging within the input VLAN are still handled entirely by hardware and not sent to the CPU.
- If the VLAN map does not fit in the hardware, all packets on that VLAN must be both filtered and forwarded by software.

Any problem in fitting the configuration into hardware is logged, but it is possible that not everyone who configures the switch can see the log messages as they occur. You can use the **show fm** privileged EXEC commands to determine if any interface configuration or VLAN configuration did not fit into hardware.

Beginning in privileged EXEC mode, follow these steps to see if a configuration fits into hardware:

	Command	Purpose
Step 1	**show fm vlan** *vlan-id*	Display feature manager information for the interface or the VLAN.
	or	
	show fm interface *interface-id*	Determine what label was used in the hardware for the interface or VLAN configuration.
Step 2	**show fm label** *name*	Display which of the configured ACL features fit into hardware.

This example shows how to display detailed feature manager information on a specified interface:

```
Switch# show fm interface gigabitethernet0/12
Input Label: 0 (default)
Output Label: 0 (default)
Priority: normal
```

This output from the **show fm label** privileged EXEC command shows a merge failure on an input access group:

```
Switch# show fm label 1
Unloaded due to merge failure or lack of space:
  InputAccessGroup
  Merge Fail:input
Input Features:
  Interfaces or VLANs: Vl1
  Priority:normal
  Vlan Map:(none)
  Access Group:131, 6788 VMRs
  Multicast Boundary:(none), 0 VMRs
Output Features:
  Interfaces or VLANs:
  Priority:low
  Bridge Group Member:no
  Vlan Map:(none)
  Access Group:(none), 0 VMRs
```

This output from the **show fm label** privileged EXEC command shows insufficient room for an input access group in the hardware:

```
Switch# show fm label 1
Unloaded due to merge failure or lack of space:
  InputAccessGroup
Input Features:
  Interfaces or VLANs: Vl1
  Priority:normal
  Vlan Map:(none)
  Access Group:bigone, 11 VMRs
  Multicast Boundary:(none), 0 VMRs
Output Features:
  Interfaces or VLANs:
  Priority:low
  Bridge Group Member:no
  Vlan Map:(none)
  Access Group:(none), 0 VMRs
```

This output from the **show fm label** privileged EXEC command shows not enough room for the input access group or the output access group on the label. (Note that the access groups were configured on two different interfaces. Labels are assigned independently for input and output.)

```
Switch# show fm label 1
Unloaded due to merge failure or lack of space:
  InputAccessGroup OutputAccessGroup
Input Features:
  Interfaces or VLANs: Vl1
  Priority:normal
  Vlan Map:(none)
  Access Group:bigone, 11 VMRs
  Multicast Boundary:(none), 0 VMRs
Output Features:
  Interfaces or VLANs: Vl2
  Priority:normal
  Bridge Group Member:no
  Vlan Map:(none)
  Access Group:bigtwo, 11 VMRs
```

Note When configuring ACLs on the switch, to allocate maximum hardware resources for ACLs, you can use the **sdm prefer access** global configuration command to set the Switch Database Management feature to the access template. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

Examples of Router ACLs and VLAN Maps Applied to VLANs

This section gives examples of applying router ACLs and VLAN maps to a VLAN for switched, bridged, routed, and multicast packets. Although the following illustrations show packets being forwarded to their destination, each time the packet's path crosses a line indicating a VLAN map or an ACL, it is also possible that the packet might be dropped, rather than forwarded.

ACLs and Switched Packets

Figure 19-6 shows how an ACL is applied on packets that are switched within a VLAN. Packets switched within the VLAN without being routed or forwarded by fallback bridging are only subject to the VLAN map of the input VLAN.

Figure 19-6 Applying ACLs on Switched Packets

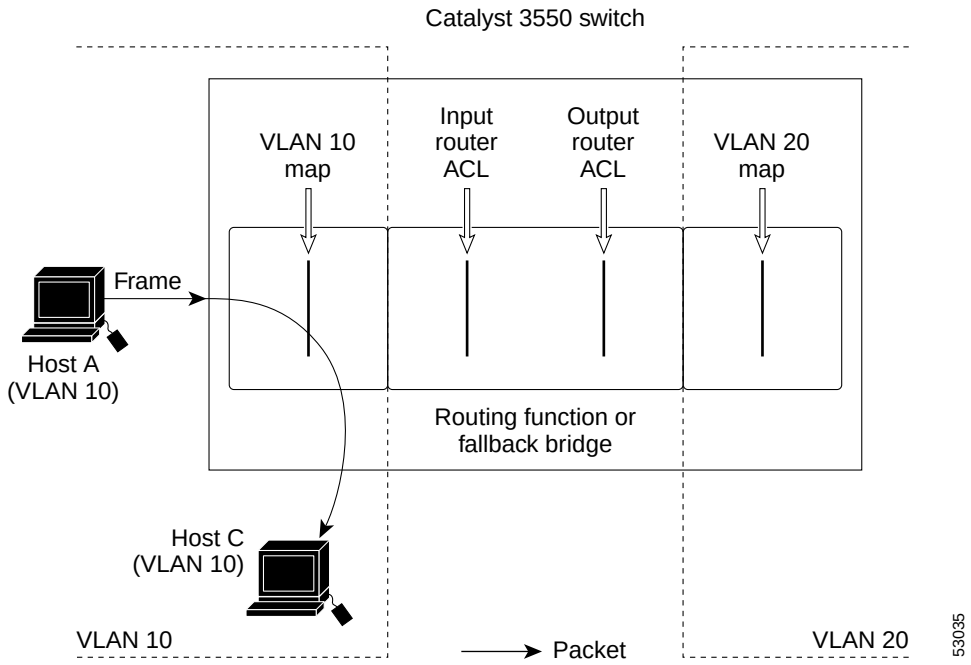

ACLs and Bridged Packets

Figure 19-7 shows how an ACL is applied on fallback-bridged packets. For bridged packets, only Layer 2 ACLs are applied to the input VLAN. Only non-IP, non-ARP packets can be fallback-bridged.

Figure 19-7 Applying ACLs on Bridged Packets

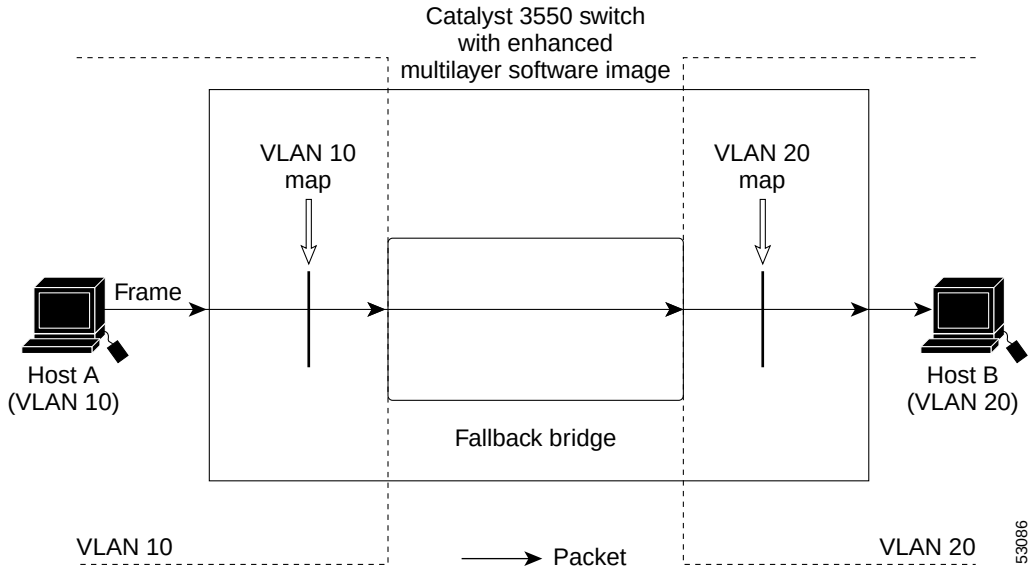

ACLs and Routed Packets

Figure 19-8 shows how ACLs are applied on routed packets. For routed packets, the ACLs are applied in this order:

1. VLAN map for input VLAN
2. Input router ACL
3. Output router ACL
4. VLAN map for output VLAN

Figure 19-8 Applying ACLs on Routed Packets

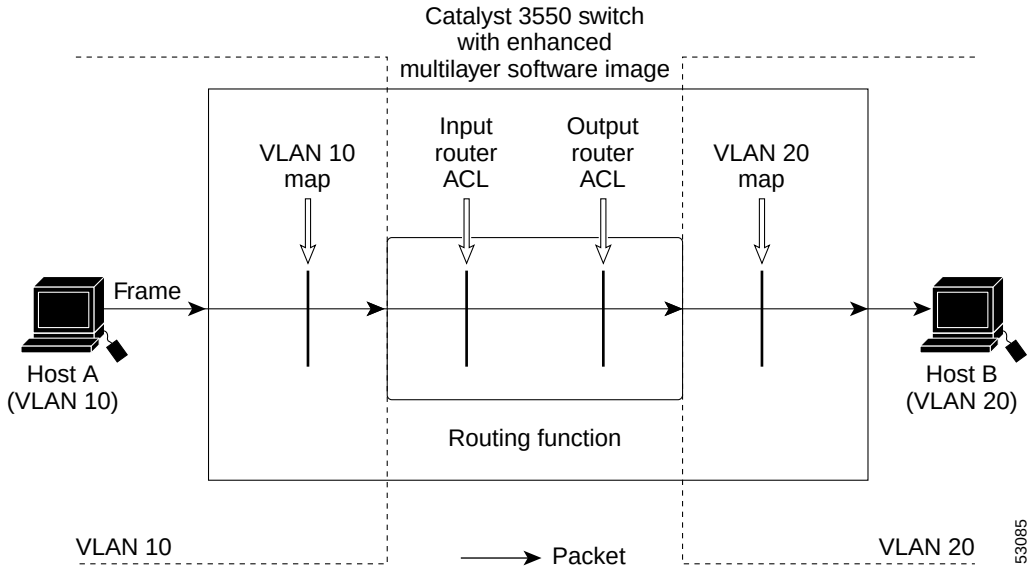

ACLs and Multicast Packets

Figure 19-9 shows how ACLs are applied on packets that are replicated for IP multicasting. A multicast packet being routed has two different kinds of filters applied: one for destinations that are other ports in the input VLAN and another for each of the destinations that are in other VLANs to which the packet has been routed. The packet might be routed to more than one output VLAN, in which case a different router output ACL and VLAN map would apply for each destination VLAN.

The final result is that the packet might be permitted in some of the output VLANs and not in others. A copy of the packet is forwarded to those destinations where it is permitted. However, if the input VLAN map (VLAN 10 map in Figure 19-9) drops the packet, no destination receives a copy of the packet.

Figure 19-9 Applying ACLs on Multicast Packets

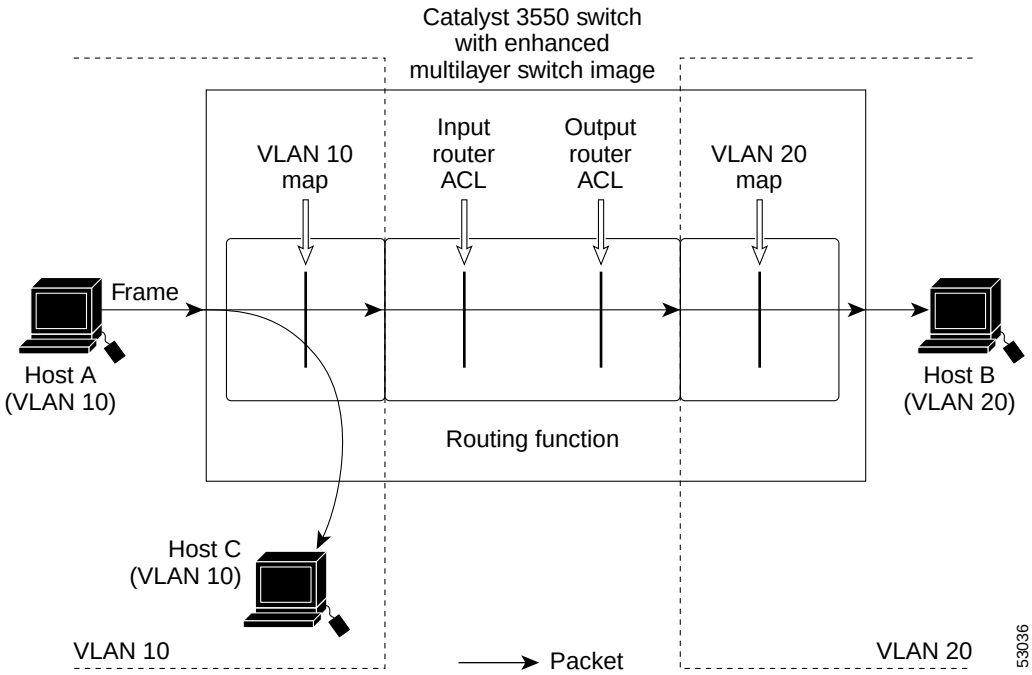

CHAPTER 20

Configuring QoS

This chapter describes how to configure quality of service (QoS) on your switch. With this feature, you can provide preferential treatment to certain traffic at the expense of others. Without QoS, the switch offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding QoS, page 20-1
- Configuring QoS, page 20-18
- Displaying QoS Information, page 20-56
- QoS Configuration Examples, page 20-56

Note When you are configuring QoS parameters for the switch, in order to allocate system resources to maximize the number of possible QoS access control entries (ACEs) allowed, you can use the **sdm prefer access** global configuration command to set the Switch Database Management feature to the access template. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

Understanding QoS

Typically, networks operate on a best-effort delivery basis, which means that all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped.

When you configure the QoS feature, you can select specific network traffic, prioritize it according to its relative importance, and use congestion-management and congestion-avoidance techniques to provide preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the DiffServ architecture, an emerging standard from the Internet Engineering Task Force (IETF). This architecture specifies that each packet is classified upon entry into the network. The classification is carried in the IP packet header, using 6 bits from the deprecated IP

type of service (TOS) field to carry the classification (*class*) information. Classification can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in Figure 20-1:

- Prioritization values in Layer 2 frames:

  Layer 2 Inter-Switch Link (ISL) frame headers have a 1-byte User field that carries an IEEE 802.1p class of service (CoS) value in the three least-significant bits. On interfaces configured as Layer 2 ISL trunks, all traffic is in ISL frames.

  Layer 2 802.1Q frame headers have a 2-byte Tag Control Information field that carries the CoS value in the three most-significant bits, which are called the User Priority bits. On interfaces configured as Layer 2 802.1Q trunks, all traffic is in 802.1Q frames except for traffic in the native VLAN.

  Other frame types cannot carry Layer 2 CoS values.

  Layer 2 CoS values range from 0 for low priority to 7 for high priority.

- Prioritization bits in Layer 3 packets:

  Layer 3 IP packets can carry either an IP precedence value or a Differentiated Services Code Point (DSCP) value. QoS supports the use of either value because DSCP values are backward-compatible with IP precedence values.

  IP precedence values range from 0 to 7.

  DSCP values range from 0 to 63.

Figure 20-1 QoS Classification Layers in Frames and Packets

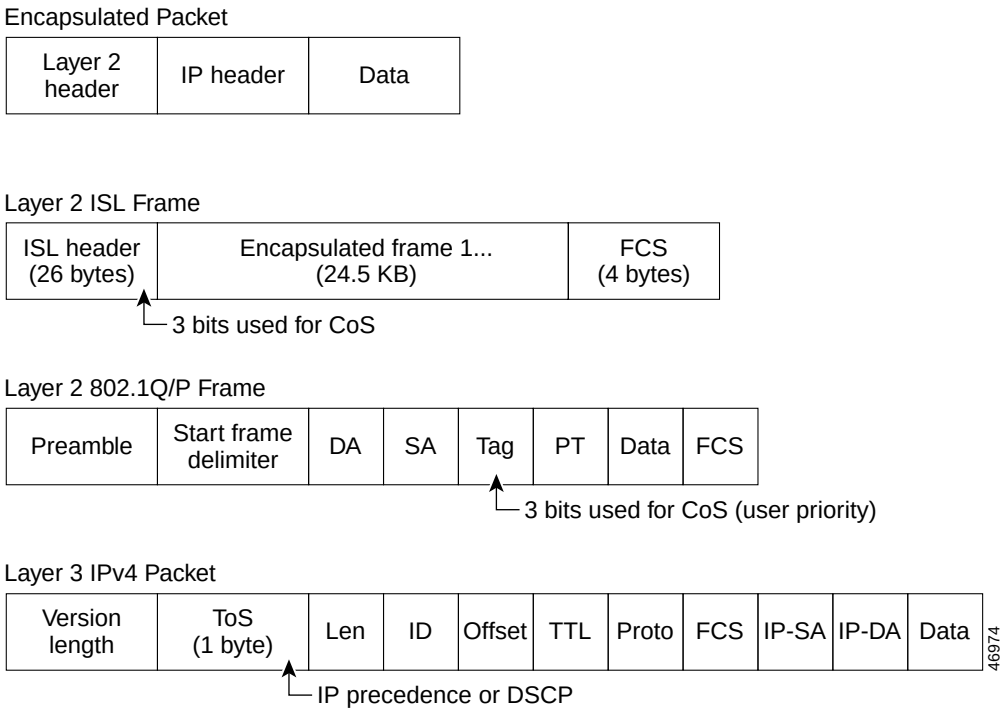

Note Layer 3 IPv6 packets are treated as non-IP packets and are bridged by the switch.

All switches and routers that access the Internet rely on the class information to provide the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by switches or routers along the way, based on a configured policy, detailed examination of the packet, or both. Detailed examination of the packet is expected to happen closer to the edge of the network so that the core switches and routers are not overloaded.

Switches and routers along the path can use the class information to limit the amount of resources allocated per traffic class. The behavior of an individual device when handling traffic in the DiffServ architecture is called per-hop behavior. If all devices along a path provide a consistent per-hop behavior, you can construct an end-to-end QoS solution.

Implementing QoS in your network can be a simple or complex task and depends on the QoS features offered by your internetworking devices, the traffic types and patterns in your network, and the granularity of control that you need over incoming and outgoing traffic.

Basic QoS Model

Figure 20-2 shows the basic QoS model. Actions at the ingress interface include classifying traffic, policing, and marking:

- Classifying distinguishes one kind of traffic from another. The process generates an internal DSCP for a packet, which identifies all the future QoS actions to be performed on this packet. For more information, see the "Classification" section on page 20-4.
- Policing determines whether a packet is in or out of profile by comparing the internal DSCP to the configured policer, and the policer limits the bandwidth consumed by a flow of traffic. The result of this determination is passed to the marker. For more information, see the "Policing and Marking" section on page 20-8.
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and decides what to do with the packet (pass through a packet without modification, mark down the DSCP value in the packet, or drop the packet). For more information, see the "Policing and Marking" section on page 20-8.

Actions at the egress interface include queueing and scheduling:

- Queueing evaluates the internal DSCP and determines which of the four egress queues in which to place the packet. The DSCP value is mapped to a CoS value, which selects one of the queues.
- Scheduling services the four egress queues based on their configured weighted round robin (WRR) weights and thresholds. One of the queues can be the expedite queue, which is serviced until empty before the other queues are serviced. Congestion avoidance techniques include tail drop and Weighted Random Early Detection (WRED) on Gigabit-capable Ethernet ports and tail drop (with only one threshold) on 10/100 Ethernet ports.

Note Policing and marking can also occur on egress interfaces.

Figure 20-2 Basic QoS Model

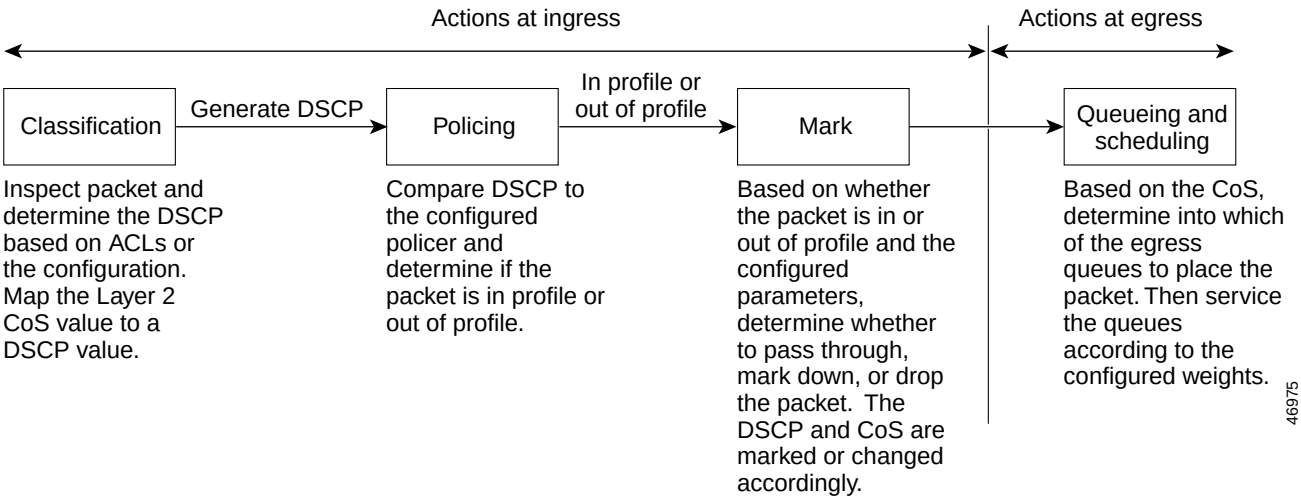
Actions at ingress
Actions at egress
Classification
Generate DSCP
Policing
In profile or out of profile
Mark
Queueing and scheduling
Inspect packet and determine the DSCP based on ACLs or the configuration. Map the Layer 2 CoS value to a DSCP value.
Compare DSCP to the configured policer and determine if the packet is in profile or out of profile.
Based on whether the packet is in or out of profile and the configured parameters, determine whether to pass through, mark down, or drop the packet. The DSCP and CoS are marked or changed accordingly.
Based on the CoS, determine into which of the egress queues to place the packet. Then service the queues according to the configured weights.
46975

Classification

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet. Classification is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled, so no classification occurs.

Note Classification occurs only on a physical interface basis. No support exists for classifying packets at the VLAN or the switch virtual interface level.

You specify which fields in the frame or packet that you want to use to classify incoming traffic.

For non-IP traffic, you have these classification options as shown in Figure 20-3:

- Use the port default. If the frame does not contain a CoS value, assign the default port CoS value to the incoming frame. Then use the configurable CoS-to-DSCP map to generate the internal DSCP value.
- Trust the CoS value in the incoming frame (configure the port to trust CoS). Then use the configurable CoS-to-DSCP map to generate the internal DSCP value. Layer 2 ISL frame headers carry the CoS value in the three least-significant bits of the 1-byte User field. Layer 2 802.1Q frame headers carry the CoS value in the three most-significant bits of the Tag Control Information field. CoS values range from 0 for low priority to 7 for high priority.

  The trust DSCP and trust IP precedence configurations are meaningless for non-IP traffic. If you configure a port with either of these options and non-IP traffic is received, the switch assigns the default port CoS value and generates the internal DSCP from the CoS-to-DSCP map.
- Perform the classification based on the configured Layer 2 MAC access control list (ACL), which can examine the MAC source address, the MAC destination address, and the Ethertype field. If no ACL is configured, the packet is assigned the default DSCP of 0, which means best-effort traffic; otherwise, the policy map specifies the DSCP to assign to the incoming frame.

For IP traffic, you have these classification options as shown in Figure 20-3:

- Trust the IP DSCP in the incoming packet (configure the port to trust DSCP), and assign the same DSCP to the packet for internal use. The IETF defines the 6 most-significant bits of the 1-byte Type of Service (ToS) field as the DSCP. The priority represented by a particular DSCP value is configurable. DSCP values range from 0 to 63.

  For ports that are on the boundary between two QoS administrative domains, you can modify the DSCP to another value by using the configurable DSCP-to-DSCP-mutation map.

- Trust the IP precedence in the incoming packet (configure the port to trust IP precedence), and generate a DSCP by using the configurable IP-precedence-to-DSCP map. The IP version 4 specification defines the three most-significant bits of the 1-byte ToS field as the IP precedence. IP precedence values range from 0 for low priority to 7 for high priority.
- Trust the CoS value (if present) in the incoming packet, and generate the DSCP by using the CoS-to-DSCP map.
- Perform the classification based on a configured IP standard or an extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned the default DSCP of 0, which means best-effort traffic; otherwise, the policy map specifies the DSCP to assign to the incoming frame.

For information on the maps described in this section, see the "Mapping Tables" section on page 20-11. For configuration information on port trust states, see the "Configuring Classification Using Port Trust States" section on page 20-21.

Figure 20-3 Classification Flowchart

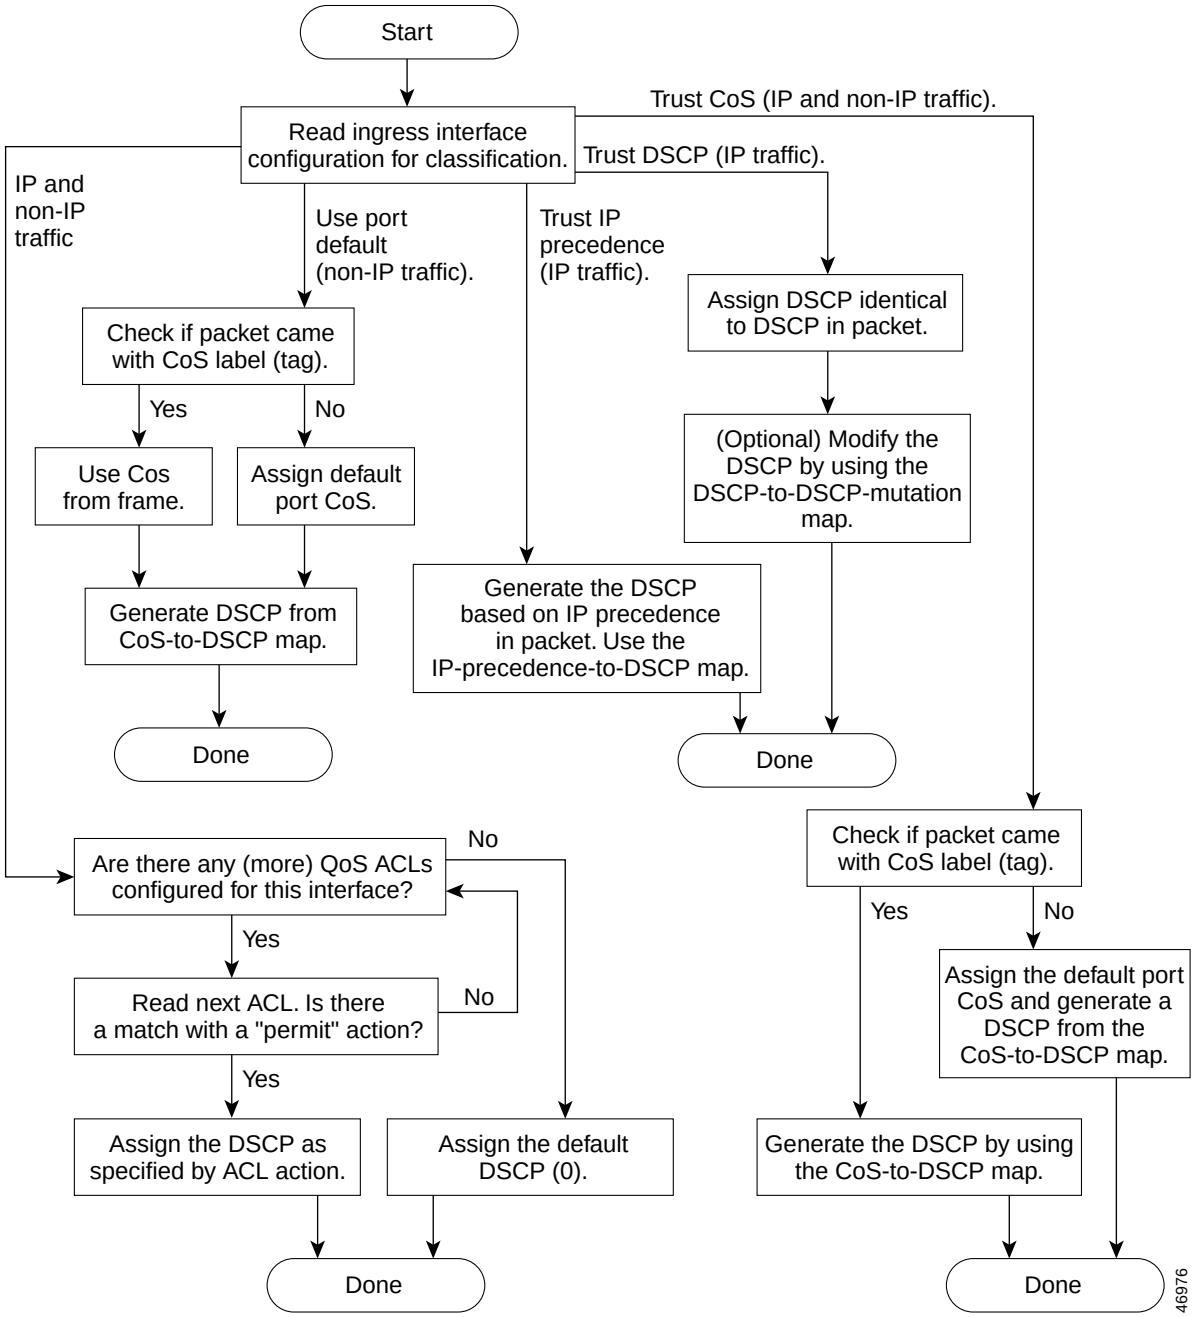
Start
Read ingress interface configuration for classification.
Trust CoS (IP and non-IP traffic).
Trust DSCP (IP traffic).
IP and non-IP traffic
Use port default (non-IP traffic).
Trust IP precedence (IP traffic).
Check if packet came with CoS label (tag).
Yes
No
Use Cos from frame.
Assign default port CoS.
Generate DSCP from CoS-to-DSCP map.
Done
Assign DSCP identical to DSCP in packet.
(Optional) Modify the DSCP by using the DSCP-to-DSCP-mutation map.
Generate the DSCP based on IP precedence in packet. Use the IP-precedence-to-DSCP map.
Done
Check if packet came with CoS label (tag).
Yes
No
Assign the default port CoS and generate a DSCP from the CoS-to-DSCP map.
Generate the DSCP by using the CoS-to-DSCP map.
Done
Are there any (more) QoS ACLs configured for this interface?
No
Yes
Read next ACL. Is there a match with a "permit" action?
No
Yes
Assign the DSCP as specified by ACL action.
Assign the default DSCP (0).
Done
46976

Classification Based on QoS ACLs

You can use IP standard, IP extended, and Layer 2 MAC ACLs to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If a match with a deny action is encountered, the ACL being processed is skipped, and the next ACL is processed.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet, and the switch offers best-effort service to the packet.
- If multiple ACLs are configured on an interface, the lookup stops after the packet matches the first ACL with a permit action, and QoS processing begins.

Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify non-IP traffic by using the **mac access-list extended** global configuration command. For configuration information, see the "Configuring a QoS Policy" section on page 20-26.

Classification Based on Class Maps and Policy Maps

A class map is a mechanism that you use to isolate and name a specific traffic flow (or class) from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it; the criteria can include matching the access group defined by the ACL or matching a specific list of DSCP or IP precedence values. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command; you should use the **class-map** command when the map is shared among many ports. When you enter the **class-map** command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class**, **trust**, or **set** policy-map configuration and policy-map class configuration commands. To make the policy map effective, you attach it to an interface by using the **service-policy** interface configuration command.

The policy map can also contain commands that define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded. For more information, see the "Policing and Marking" section on page 20-8.

A policy map also has these characteristics:

- A policy map can contain multiple class statements.
- A separate policy-map class can exist for each type of traffic received through an interface.
- The policy-map trust state and an interface trust state are mutually exclusive, and whichever is configured last takes affect.

For configuration information, see the "Configuring a QoS Policy" section on page 20-26.

Policing and Marking

After a packet is classified and has an internal DSCP value assigned to it, the policing and marking process can begin as shown in Figure 20-4.

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer specifies the action to take for packets that are in or out of profile. These actions, carried out by the marker, include passing through the packet without modification, dropping the packet, or marking down the packet with a new DSCP value that is obtained from the configurable policed-DSCP map. For information on the policed-DSCP map, see the "Mapping Tables" section on page 20-11.

You can create these types of policers:

- Individual

  QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **police** policy-map configuration command.

- Aggregate

  QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all matched traffic flows. You configure this type of policer by specifying the aggregate policer name within a policy map by using the **police aggregate** policy-map configuration command. You specify the bandwidth limits of the policer by using the **mls qos aggregate-policer** global configuration command. In this way, the aggregate policer is shared by multiple classes of traffic within a policy map.

Policing uses a token bucket algorithm. As each frame is received by the switch, a token is added to the bucket. The bucket has a hole in it and leaks at a rate that you specify as the average traffic rate in bits per second. Each time a token is added to the bucket, the switch performs a check to determine if there is enough room in the bucket. If there is not enough room, the packet is marked as nonconforming, and the specified policer action is taken (dropped or marked down).

How quickly the bucket fills is a function of the bucket depth (burst-byte), the rate at which the tokens are removed (rate-bps), and the duration of the burst above the average rate. The size of the bucket imposes an upper limit on the burst length and determines the number of frames that can be transmitted back-to-back. If the burst is short, the bucket does not overflow, and no action is taken against the traffic flow. However, if a burst is long and at a higher rate, the bucket overflows and the policing actions are taken against the frames in that burst.

You configure the bucket depth (the maximum burst that is tolerated before the bucket overflows) by using the *burst-byte* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command. You configure how quickly (the average rate) the tokens are removed from the bucket by using the *rate-bps* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command.

When configuring policing and policers, keep these items in mind:

- By default, no policers are configured.
- Policers can be configured only on a physical port. There is no support for policing at a VLAN or switch virtual interface level.
- Only one policer can be applied to a packet per direction.
- Only the average rate and committed burst parameters are configurable.
- Policing can occur on ingress and egress interfaces:
  - 128 policers are supported on ingress Gigabit-capable Ethernet ports.
  - 8 policers are supported on ingress 10/100 Ethernet ports.
  - 8 policers are supported on all egress ports.
  - Ingress policers can be individual or aggregate.
- On an interface configured for QoS, all traffic received through the interface is classified, policed, and marked according to the policy map attached to the interface. On a trunk interface configured for QoS, traffic in *all* VLANs received through the interface is classified, policed, and marked according to the policy map attached to the interface.

After you configure the policy map and policing actions, attach the policy to an ingress or egress interface by using the **service-policy** interface configuration command. For configuration information, see the "Classifying, Policing, and Marking Traffic by Using Policy Maps" section on page 20-32 and the "Classifying, Policing, and Marking Traffic by Using Aggregate Policers" section on page 20-37.

Figure 20-4 Policing and Marking Flowchart

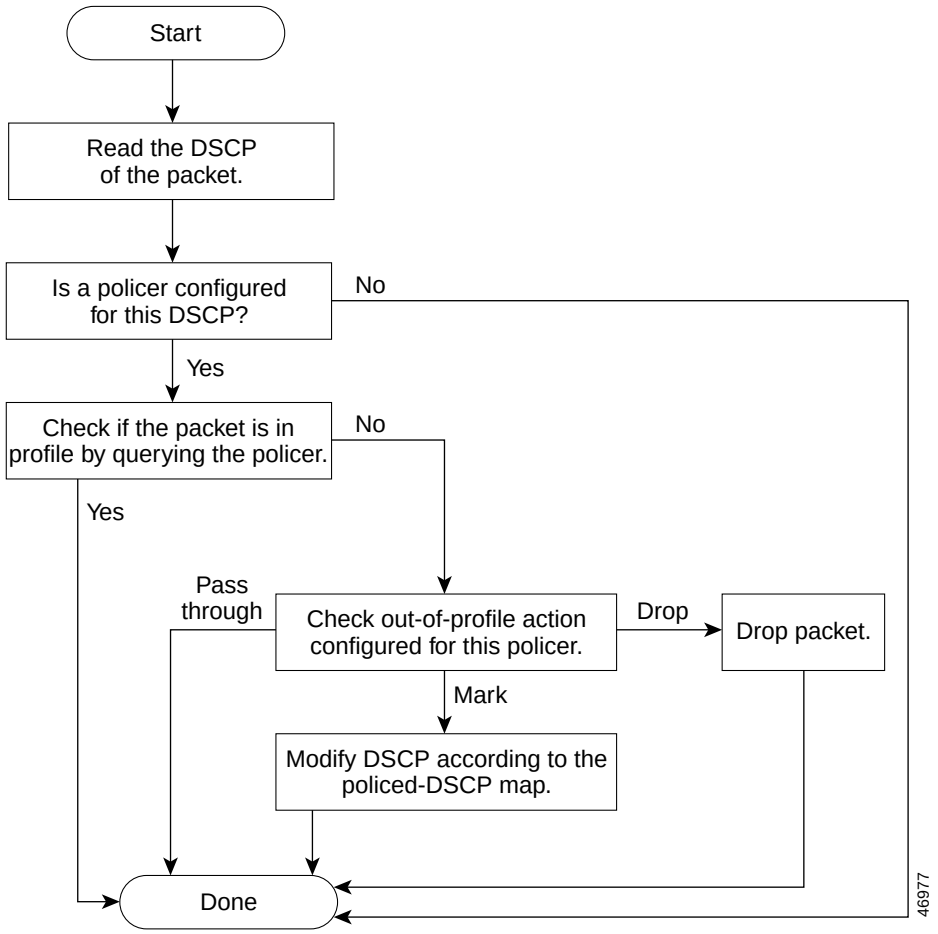

Mapping Tables

During QoS processing, the switch represents the priority of all traffic (including non-IP traffic) with an internal DSCP value:

- During classification, QoS uses configurable mapping tables to derive the internal DSCP (a 6-bit value) from received CoS or IP precedence (3-bit) values. These maps include the CoS-to-DSCP map and the IP-precedence-to-DSCP map.

  On an ingress interface configured in the DSCP-trusted state, if the DSCP values are different between the QoS domains, you can apply the configurable DSCP-to-DSCP-mutation map to the interface that is on the boundary between the two QoS domains.

- During policing, QoS can assign another DSCP value to an IP or non-IP packet (if the packet is out of profile and the policer specifies a marked down DSCP value). This configurable map is called the policed-DSCP map.

- Before the traffic reaches the scheduling stage, QoS uses the configurable DSCP-to-CoS map to derive a CoS value from the internal DSCP value. Through the CoS-to-egress-queue map, the CoS values select one of the four egress queues for output processing.

The CoS-to-DSCP, DSCP-to-CoS, and the IP-precedence-to-DSCP map have default values that might or might not be appropriate for your network.

The default DSCP-to-DSCP-mutation map and the default policed-DSCP map are null maps; they map an incoming DSCP value to the same DSCP value. The DSCP-to-DSCP-mutation map is the only map you apply to a specific Gigabit-capable Ethernet port or to a group of 10/100 Ethernet ports. All other maps apply to the entire switch.

For configuration information, see the "Configuring DSCP Maps" section on page 20-39.

Queueing and Scheduling

After a packet is policed and marked, the queueing and scheduling process begins as described in these sections:

- Queueing and Scheduling on Gigabit-Capable Ports, page 20-12
- Queueing and Scheduling on 10/100 Ethernet Ports, page 20-15

Queueing and Scheduling on Gigabit-Capable Ports

Figure 20-5 shows the queueing and scheduling flowchart for Gigabit-capable Ethernet ports.

Figure 20-5 Queueing and Scheduling Flowchart for Gigabit-Capable Ethernet Ports

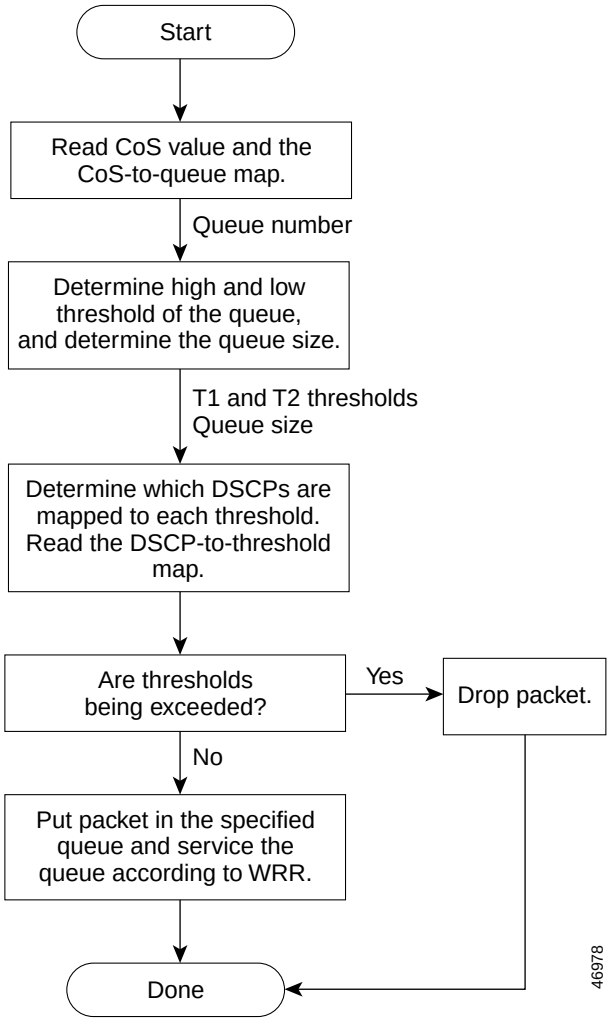
Start
Read CoS value and the CoS-to-queue map.
Queue number
Determine high and low threshold of the queue, and determine the queue size.
T1 and T2 thresholds
Queue size
Determine which DSCPs are mapped to each threshold. Read the DSCP-to-threshold map.
Are thresholds being exceeded?
Yes
Drop packet.
No
Put packet in the specified queue and service the queue according to WRR.
Done
46978

Note If the expedite queue is enabled, WRR services it until it is empty before servicing the other three queues.

During the queueing and scheduling process, the switch uses egress queues and WRR for congestion management, and tail drop or WRED algorithms for congestion avoidance on Gigabit-capable Ethernet ports.

Each Gigabit-capable Ethernet port has four egress queues, one of which can be the egress expedite queue. You can configure the buffer space allocated to each queue as a ratio of weights by using the **wrr-queue queue-limit** interface configuration command, where the relative size differences in the numbers indicates the relative differences in the queue sizes. To display the absolute value of the queue size, use the **show mls qos interface** *interface-id* **statistics** privileged EXEC command, and examine the FreeQ information.

You assign two drop thresholds to each queue, map DSCPs to the thresholds through the DSCP-to-threshold map, and enable either tail drop or WRED on the interface. The queue size, drop thresholds, tail-drop or WRED algorithm, and the DSCP-to-threshold map work together to determine when and which packets are dropped when the thresholds are exceeded. You configure the drop percentage thresholds by using either the **wrr-queue threshold** interface configuration command for tail drop or the **wrr-queue random-detect max-threshold** interface configuration command for WRED; in either case, you map DSCP values to the thresholds (DSCP-to-threshold map) by using the **wrr-queue dscp-map** interface configuration command. For more information, see the "Tail Drop" section on page 20-13 and "WRED" section on page 20-14.

The available bandwidth of the egress link is divided among the queues. You configure the queues to be serviced according to the ratio of WRR weights by using the **wrr-queue bandwidth** interface configuration command. Queues are selected by the CoS value that is mapped to an egress queue (CoS-to-egress-queue map) through the **wrr-queue cos-map** interface configuration command.

All four queues participate in the WRR unless the expedite queue is enabled, in which case, the fourth bandwidth weight is ignored and not used in the ratio calculation. The expedite queue is a strict-priority queue, and it is serviced until empty before the other queues are serviced. You enable the expedite queue by using the **priority-queue out** interface configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs into certain queues, allocate a larger queue size or service the particular queue more frequently, and adjust queue thresholds so that packets with lower priorities are dropped. For configuration information, see the "Configuring Egress Queues on Gigabit-Capable Ethernet Ports" section on page 20-44.

Tail Drop

Tail drop is the default congestion-avoidance technique on Gigabit-capable Ethernet ports. With tail drop, packets are queued until the thresholds are exceeded. Specifically, all packets with DSCPs assigned to the first threshold are dropped until the threshold is no longer exceeded. However, packets assigned to the second threshold continue to be queued and sent as long as the second threshold is not exceeded.

You can modify the two tail-drop threshold percentages assigned to the four egress queues by using the **wrr-queue threshold** interface configuration command. Each threshold value is a percentage of the total number of allocated queue descriptors for the queue. The default threshold is 100 percent for thresholds 1 and 2.

You modify the DSCP-to-threshold map to determine which DSCPs are mapped to which threshold ID by using the **wrr-queue dscp-map** interface configuration command. By default, all DSCPs are mapped to threshold 1, and when this threshold is exceeded, all the packets are dropped.

If you use tail-drop thresholds, you cannot use WRED, and vice versa. If tail drop is disabled, WRED is automatically enabled with the previous configuration (or the default if it was not previously configured).

WRED

Cisco's implementation of Random Early Detection (RED), called Weighted Random Early Detection (WRED), differs from other congestion-avoidance techniques because it attempts to anticipate and avoid congestion, rather than controlling congestion once it occurs.

WRED takes advantage of TCP congestion control to try to control the average queue size by indicating to end hosts when they should temporarily stop sending packets. By randomly dropping packets before periods of high congestion, it tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, WRED tells it to decrease its transmission rate until all the packets reach their destination, meaning that the congestion is cleared.

WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once. Thus, WRED allows the transmission line to be fully used at all times. WRED also drops more packets from large users than small. Therefore, traffic sources that generate the most traffic are more likely to be slowed down than traffic sources that generate little traffic.

You can enable WRED and configure the two threshold percentages assigned to the four egress queues on a Gigabit-capable Ethernet port by using the **wrr-queue random-detect max-threshold** interface configuration command. Each threshold percentage represents where WRED starts to randomly drop packets. After a threshold is exceeded, WRED randomly begins to drop packets assigned to this threshold. As the queue limit is approached, WRED continues to drop more and more packets. When the queue limit is reached, WRED drops all packets assigned to the threshold. By default, WRED is disabled.

You modify the DSCP-to-threshold map to determine which DSCPs are mapped to which threshold ID by using the **wrr-queue dscp-map** interface configuration command. By default, all DSCPs are mapped to threshold 1, and when this threshold is exceeded, all the packets are randomly dropped.

If you use WRED thresholds, you cannot use tail drop, and vice versa. If WRED is disabled, tail drop is automatically enabled with the previous configuration (or the default if it was not previously configured).

Queueing and Scheduling on 10/100 Ethernet Ports

Figure 20-6 shows the queueing and scheduling flowchart for 10/100 Ethernet ports.

Figure 20-6 Queueing and Scheduling Flowchart for 10/100 Ethernet Ports

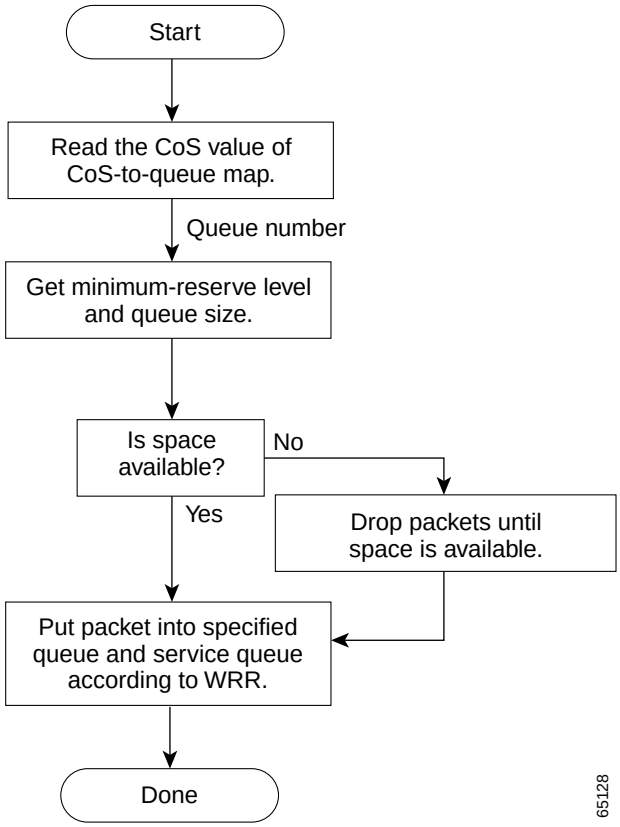

Note If the expedite queue is enabled, WRR services it until it is empty before servicing the other three queues.

During the queueing and scheduling process, the switch uses egress queues, which select the minimum-reserve level and buffer size, and WRR to provide congestion management.

Each 10/100 Ethernet port has four egress queues, one of which can be the egress expedite queue. Each queue can access one of eight minimum-reserve levels; each level has 100 packets of buffer space by default for queueing packets. When the buffer specified for the minimum-reserve level is full, packets are dropped until space is available.

Figure 20-7 is an example of the 10/100 Ethernet port queue assignments, minimum-reserve levels, and buffer sizes. The figure shows four egress queues per port, with each queue assigned to a minimum-reserve level. For example, for Fast Ethernet port 0/1, queue 1 is assigned to minimum-reserve level 1, queue 2 is assigned to minimum-reserve level 3, queue 3 is assigned to minimum-reserve level 5, and queue 4 is assigned to minimum-reserve level 7. You assign the minimum-reserve level to a queue by using the **wrr-queue min-reserve** interface configuration command.

Each minimum-reserve level is configured with a buffer size. As shown in the figure, queue 4 of Fast Ethernet port 0/1 has a buffer size of 70 packets, queue 4 of Fast Ethernet port 0/2 has a buffer size of 80 packets, queue 4 of Fast Ethernet port 0/3 has a buffer size of 40 packets, and Fast Ethernet port 0/4 has a buffer size of 80 packets. You configure the buffer size by using the **mls qos min-reserve** global configuration command.

Figure 20-7 10/100 Ethernet Port Queue Assignment, Minimum-Reserve Levels, and Buffer Size

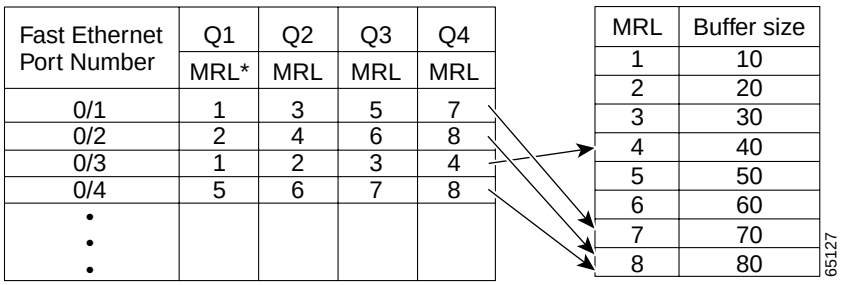

* MRL = Minimum-reserve level

The available bandwidth of the egress link is divided among the queues. You configure the queues to be serviced according to the ratio of WRR weights by using the **wrr-queue bandwidth** interface configuration command. Queues are selected by the CoS value that is mapped to an egress queue (CoS-to-egress-queue map) through the **wrr-queue cos-map** interface configuration command.

All four queues participate in the WRR unless the egress expedite queue is enabled, in which case, the fourth bandwidth weight is ignored and not used in the ratio calculation. The expedite queue is a strict-priority queue, and it is serviced until empty before the other queues are serviced. You enable the expedite queue by using the **priority-queue out** interface configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs into certain queues, allocate a larger minimum-reserve buffer size, and service a particular queue more frequently. For configuration information, see the "Configuring Egress Queues on 10/100 Ethernet Ports" section on page 20-51.

Packet Modification

A packet is classified, policed, and queued to provide QoS. Packet modifications can occur during this process:

- For IP packets, classification involves assigning a DSCP to the packet. However, the packet is not modified at this stage; only an indication of the assigned DSCP is carried along. The reason for this is that QoS classification and ACL lookup occur in parallel, and it is possible that the ACL specifies that the packet should be denied and logged. In this situation, the packet is forwarded with its original DSCP to the CPU, where it is again processed through ACL software. However, route lookup is performed based on classified DSCPs.
- For non-IP packets, classification involves assigning an internal DSCP to the packet, but because there is no DSCP in the non-IP packet, no overwrite occurs. Instead, the internal DSCP is translated to the CoS and is used both for queueing and scheduling decisions and for writing the CoS priority value in the tag if the packet is being sent on either an ISL or 802.1Q trunk port. Because the CoS priority is written in the tag, Catalyst 3500 series XL switches that use the 802.1P priority can interoperate with the QoS implementation on the Catalyst 3550 switches.
- During policing, IP and non-IP packets can have another DSCP assigned to them (if they are out of profile and the policer specifies a markdown DSCP). Once again, the DSCP in the packet is not modified, but an indication of the marked-down value is carried along. For IP packets, the packet modification occurs at a later stage; for non-IP packets the DSCP is converted to CoS and used for queueing and scheduling decisions.

Configuring QoS

Before configuring QoS, you must have a thorough understanding of these items:

- The types of applications used and the traffic patterns on your network.
- Traffic characteristics and needs of your network. Is the traffic bursty? Do you need to reserve bandwidth for voice and video streams?
- Bandwidth requirements and speed of the network.
- Location of congestion points in the network.

This section describes how to configure QoS on your switch. It contains this configuration information:

- Default QoS Configuration, page 20-18
- Enabling QoS Globally, page 20-21
- Configuring Classification Using Port Trust States, page 20-21
- Configuring a QoS Policy, page 20-26
- Configuring DSCP Maps, page 20-39
- Configuring Egress Queues on Gigabit-Capable Ethernet Ports, page 20-44
- Configuring Egress Queues on 10/100 Ethernet Ports, page 20-51

Default QoS Configuration

Table 20-1 shows the default QoS configuration when QoS is disabled.

Table 20-1 Default QoS Parameters when QoS is Disabled

Port Type	QoS State	Egress traffic (DSCP and CoS Value)	Queue	Queue Weights	Tail-drop Thresholds	CoS Mapping to Queue
Gigabit-capable Ethernet ports	Disabled	Pass through.	All of the queue RAM is allocated to queue 1 (no expedite queue).	–	100%, 100% WRED is disabled.	All CoS values map to queue 1.
10/100 Ethernet ports	Disabled	Pass through.	Each of the eight minimum-reserve levels have a buffer size of 100 packets. The queue selects the level.	–	–	All CoS values map to queue 1.

When QoS is disabled, there is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed).

Table 20-2 shows the default QoS parameters without any further configuration when QoS is enabled.

Table 20-2 Default QoS Parameters when QoS is Enabled

Port Type	QoS State	Egress traffic (DSCP and CoS Value)	Queue	Queue Weights	Tail-drop Thresholds	CoS Mapping to Queue
Gigabit-capable Ethernet ports	Enabled (no policing)	DSCP=0 CoS=0 (0 means best-effort delivery.)	Four queues are available (no expedite queue).	Each queue has the same weight.	100%, 100% WRED is disabled.	0, 1: queue 1 2, 3: queue 2 4, 5: queue 3 6, 7: queue 4
10/100 Ethernet ports	Enabled (no policing)	DSCP=0 CoS=0 (0 means best-effort delivery.)	Each of the eight minimum-reserve levels have a buffer size of 100 packets. The queue selects the level.	Each queue has the same weight.	–	0, 1: queue 1 2, 3: queue 2 4, 5: queue 3 6, 7: queue 4

The default port CoS value is 0.

The default port trust state on all ports is untrusted.

No policy maps are configured.

No policers are configured.

The default CoS-to-DSCP map is shown in Table 20-3 on page 20-39.

The default IP-precedence-to-DSCP map is shown in Table 20-4 on page 20-40.

The default DSCP-to-CoS map is shown in Table 20-5 on page 20-42.

The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

The default policed-DSCP map is a null map, which maps an incoming DSCP value to the same DSCP value (no markdown).

The default DSCP-to-switch-priority map maps DSCPs 0 to 15 to priority 0, DSCPs 16 to 31 to priority 1, DSCPs 32 to 47 to priority 2, and DSCPs 48 to 63 to priority 3.

Configuration Guidelines

Before beginning the QoS configuration, you should be aware of this information:

- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.
- You configure QoS only on physical ports; there is no support for it on the VLAN or switch virtual interface level.
- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort. IP fragments are denoted by fields in the IP header.
- You can match IP options against configured IP extended ACLs to enforce QoS. These packets are sent to the CPU and processed by software. IP options are denoted by fields in the IP header.
- Control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) received by the switch are subject to all ingress QoS processing.
- You must disable the IEEE 802.3X flowcontrol on all ports before enabling QoS on the switch. To disable it, use the **flowcontrol receive off** and **flowcontrol send off** interface configuration commands.
- Only one ACL per class map and only one **match** class-map configuration command per class map are supported. The ACL can have multiple access control entries, which are commands that match fields against the contents of the packet.
- Use only the **match ip dscp** *dscp-list* class-map configuration command in a policy map that is attached to an egress interface.
- Policy maps with ACL classification in the egress direction are not supported and cannot be attached to an interface by using the **service-policy output** *policy-map-name* interface configuration command. Policy maps containing **set** or **trust** policy-map class configuration commands cannot be attached to an egress interface; instead, you can use the **police** policy-map class configuration command to mark down (reduce) the DSCP value at the egress interface.
- You can create an aggregate policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps or interfaces.

Enabling QoS Globally

By default, QoS is disabled on the switch, which means that the switch offers best-effort service to each packet regardless of the packet contents or size. All CoS values map to egress queue 1 with both tail-drop thresholds set to 100 percent of the total queue size for Gigabit-capable Ethernet ports. On 10/100 Ethernet ports, all CoS values map to egress queue 1, which uses minimum-reserve level 1 and can hold up to 100 packets. When the buffer is full, packets are dropped.

Beginning in privileged EXEC mode, follow these steps to enable QoS:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface range** *port-range*	Enter interface configuration mode, and execute a command on multiple interfaces. You can define up to five interface ranges with a single command, with each range separated by a comma. All interfaces in a range must be the same type; that is, all Fast Ethernet ports or all Gigabit Ethernet ports.
Step 3	**flowcontrol receive off** **flowcontrol send off**	Disable flowcontrol on all interfaces.
Step 4	**exit**	Return to global configuration mode.
Step 5	**mls qos**	Enable QoS globally.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show mls qos**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

After QoS is enabled, the default settings are as shown in Table 20-1 on page 20-18.

To disable QoS, use the **no mls qos** global configuration command.

Configuring Classification Using Port Trust States

This section describes how to classify incoming traffic by using port trust states. It contains this configuration information:

- Configuring the Trust State on Ports within the QoS Domain, page 20-22
- Configuring the CoS Value for an Interface, page 20-24
- Configuring the DSCP Trust State on a Port Bordering Another QoS Domain, page 20-25

Configuring the Trust State on Ports within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain. Figure 20-8 shows a sample network topology.

Figure 20-8 Port Trusted States within the QoS Domain

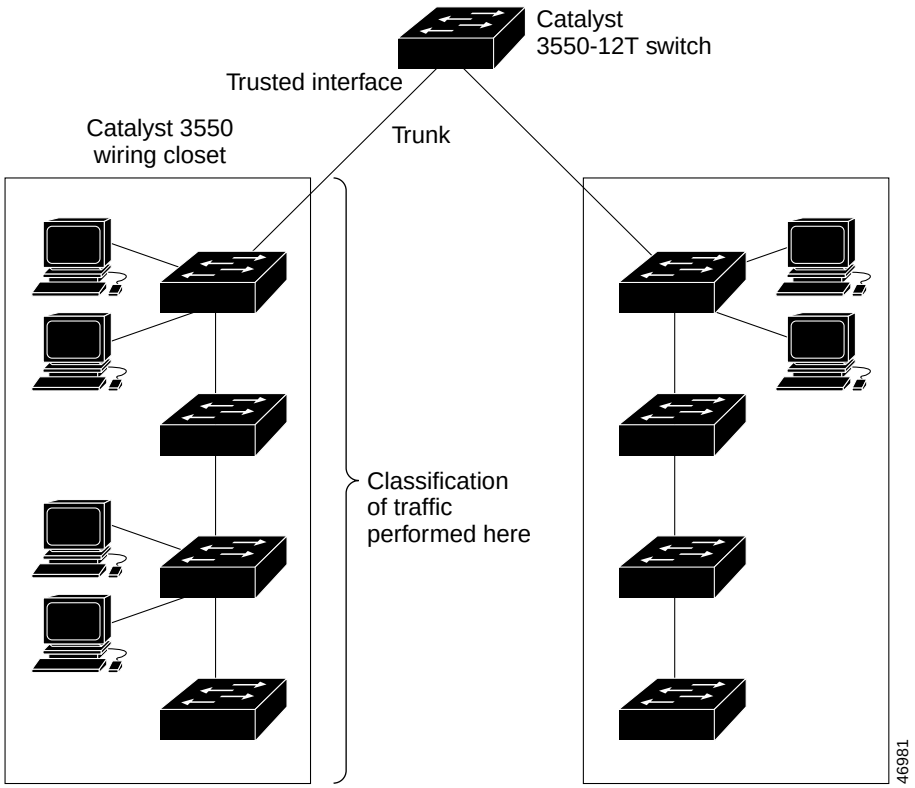

Beginning in privileged EXEC mode, follow these steps to configure the port to trust the classification of the traffic that it receives:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS globally.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 4	**mls qos trust** {**cos** | **dscp** | **ip-precedence**}	Configure the port trust state. By default, the port is not trusted. The keywords have these meanings: • **cos**—Classifies ingress packets with the packet CoS values. For untagged packets, the port default CoS value is used. The default port CoS value is 0. • **dscp**—Classifies ingress packets with packet DSCP values. For non-IP packets, the packet CoS value is used if the packet is tagged; for untagged packets, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map. • **ip-precedence**—Classifies ingress packets with the packet IP-precedence values. For non-IP packets, the packet CoS value is used if the packet is tagged; for untagged packets, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map. Use the **cos** keyword setting if your network is composed of Ethernet LANs, Catalyst 3500 XL and 2900 XL switches, and has no more than two types of traffic. Recall that on Catalyst 3500 XL and 2900 XL switches, CoS configures each transmit port with a normal-priority transmit queue and a high-priority transmit queue. Use the **dscp** or **ip-precedence** keyword if your network is not composed of only Ethernet LANs and if you are familiar with sophisticated QoS features and implementations.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return a port to its untrusted state, use the **no mls qos trust** interface configuration command.

For information on how to change the default CoS value, see the "Configuring the CoS Value for an Interface" section on page 20-24. For information on how to configure the CoS-to-DSCP map, see the "Configuring the CoS-to-DSCP Map" section on page 20-39.

Configuring the CoS Value for an Interface

QoS assigns the CoS value specified with the **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports.

Beginning in privileged EXEC mode, follow these steps to define the default CoS value of a port or to assign the default CoS to all incoming packets on the port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS globally.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 4	**mls qos cos** {*default-cos* | **override**}	Configure the default CoS value for the port. • For *default-cos*, specify a default CoS value to be assigned to a port. If the port is CoS trusted and packets are untagged, the default CoS value becomes the CoS value for the packet. The CoS range is 0 to 7. The default is 0. • Use the **override** keyword to override the previously configured trust state of the incoming packets and to apply the default port CoS value to all incoming packets. By default, CoS override is disabled. Use the **override** keyword when all incoming packets on certain ports deserve higher or lower priority than packets entering from other ports. Even if a port was previously set to trust DSCP, CoS, or IP precedence, this command overrides the previously configured trust state, and all the incoming CoS values are assigned the default CoS value configured with this command. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the ingress port.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos cos** {*default-cos* | **override**} interface configuration command.

Configuring the DSCP Trust State on a Port Bordering Another QoS Domain

If you are administering two separate QoS domains between which you want to implement QoS features for IP traffic, you can configure the switch ports bordering the domains to a DSCP-trusted state as shown in Figure 20-9. Then the receiving port accepts the DSCP-trusted value and avoids the classification stage of QoS. If the two domains use different DSCP values, you can configure the DSCP-to-DSCP-mutation map to translate a set of DSCP values to match the definition in the other domain.

Figure 20-9 DSCP-Trusted State on a Port Bordering Another QoS Domain

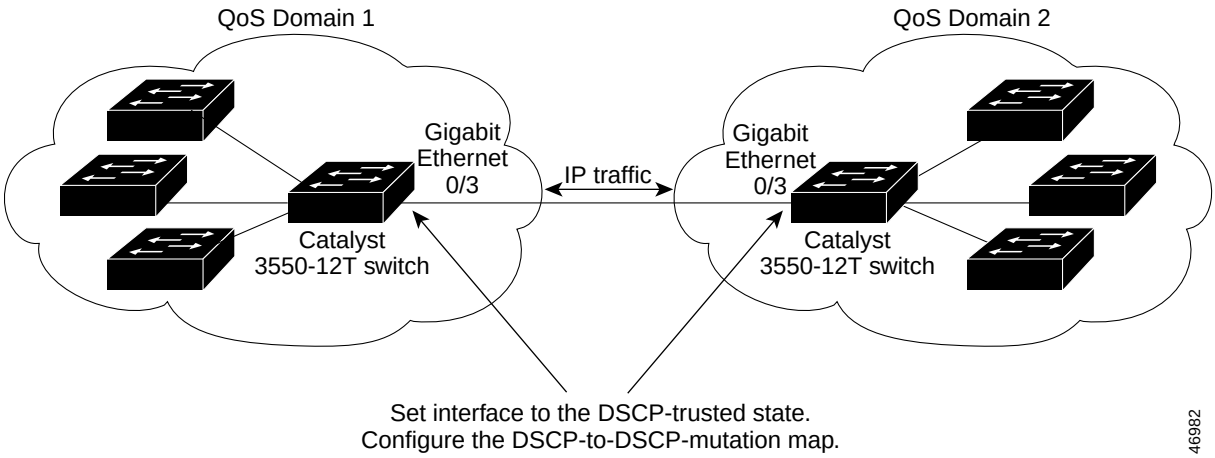

Beginning in privileged EXEC mode, follow these steps to configure the DSCP-trusted state on a port and modify the DSCP-to-DSCP-mutation map. To ensure a consistent mapping strategy across both QoS domains, you must perform this procedure on the ports in both domains:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**mls qos map dscp-mutation** *dscp-mutation-name in-dscp* **to** *out-dscp*	Modify the DSCP-to-DSCP-mutation map. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value. • For *dscp-mutation-name*, enter the mutation map name. You can create more than one map by specifying a new name. • For *in-dscp*, enter up to eight DSCP values separated by spaces. Then enter the **to** keyword. • For *out-dscp*, enter up to eight DSCP values separated by spaces. The DSCP range is 0 to 63.
Step 4	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 5	**mls qos trust dscp**	Configure the ingress port as a DSCP-trusted port.

	Command	Purpose
Step 6	**mls qos dscp-mutation** *dscp-mutation-name*	Apply the map to the specified ingress DSCP-trusted port. You can apply the map to different Gigabit-capable Ethernet ports. However, on 10/100 Ethernet ports, you can attach only one DSCP-to-DSCP-mutation map to a group of twelve ports. For example, Fast Ethernet ports 0/1 to 0/12 are a group, Fast Ethernet ports 0/13 to 0/24 are a group, Gigabit Ethernet 0/1 is a group, and Gigabit Ethernet 0/2 is a group. When applying a mutation map to any port in a group, all ports in the same group are automatically configured with the same map.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show mls qos maps dscp-mutation**	Verify your entries.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return a port to its non-trusted state, use the **no mls qos trust** interface configuration command. To return to the default DSCP-to-DSCP-mutation map values, use the **no mls qos map dscp-mutation** *dscp-mutation-map-name* global configuration command.

This example shows how to configure Gigabit Ethernet port 0/2 to the DSCP-trusted state and to modify the DSCP-to-DSCP-mutation map (named *gi0/2-mutation*) so that incoming DSCP values 10 to 13 are mapped to DSCP values 30 to 33:

```
Switch# configure terminal
Switch(config)# mls qos map dscp-mutation gi0/2-mutation 10 11 12 13 to 30 31 32 33
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation gi0/2-mutation
Switch(config-if)# end
```

Configuring a QoS Policy

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to interfaces.

For background information, see the "Classification" section on page 20-4 and the "Policing and Marking" section on page 20-8.

This section contains this configuration information:

- Classifying Traffic by Using ACLs, page 20-27
- Classifying Traffic by Using Class Maps, page 20-30
- Classifying, Policing, and Marking Traffic by Using Policy Maps, page 20-32
- Classifying, Policing, and Marking Traffic by Using Aggregate Policers, page 20-37

Classifying Traffic by Using ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify non-IP traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for IP traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create an IP standard ACL, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number. The range is 1 to 99 and 1300 to 1999. • Use the **permit** keyword to permit a certain type of traffic if the conditions are matched. Use the **deny** keyword to deny a certain type of traffic if conditions are matched. • For *source*, enter the network or host from which the packet is being sent. You can use the **any** keyword as an abbreviation for 0.0.0.0 255.255.255.255. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to allow access for only those hosts on the three specified networks. The wildcard bits apply to the host portions of the network addresses. Any host with a source address that does not match the access list statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 128.88.0.0 0.0.255.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
! (Note: all other access implicitly denied)
```

Beginning in privileged EXEC mode, follow these steps to create an IP extended ACL for IP traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	Create an IP extended ACL, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number. The range is 100 to 199 and 2000 to 2699. • Use the **permit** keyword to permit a certain type of traffic if the conditions are matched. Use the **deny** keyword to deny a certain type of traffic if conditions are matched. • For *protocol*, enter the name or number of an IP protocol. Use the question mark (?) to see a list of available protocol keywords. • For *source*, enter the network or host from which the packet is being sent. You specify this by using dotted decimal notation, by using the **any** keyword as an abbreviation for *source* 0.0.0.0 *source-wildcard* 255.255.255.255, or by using the **host** keyword for *source* 0.0.0.0. • For *source-wildcard*, enter the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the wildcard by using dotted decimal notation, by using the **any** keyword as an abbreviation for *source* 0.0.0.0 *source-wildcard* 255.255.255.255, or by using the **host** keyword for *source* 0.0.0.0. • For *destination*, enter the network or host to which the packet is being sent. You have the same options for specifying the *destination* and *destination-wildcard* as those described by *source* and *source-wildcard*. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show access-lists**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to create an ACL that permits IP traffic from any source to any destination that has the DSCP value set to 32:

```
Switch(config)# access-list 100 permit ip any any dscp 32
```

This example shows how to create an ACL that permits IP traffic from a source host at 10.1.1.1 to a destination host at 10.1.1.2 with a precedence value of 5:

```
Switch(config)# access-list 100 permit ip host 10.1.1.1 host 10.1.1.2 precedence 5
```

This example shows how to create an ACL that permits PIM traffic from any source to a destination group address of 224.0.0.2 with a DSCP set to 32:

```
Switch(config)# access-list 102 permit pim any 224.0.0.2 dscp 32
```

Beginning in privileged EXEC mode, follow these steps to create a Layer 2 MAC ACL for non-IP traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**mac access-list extended** *name*	Create a Layer 2 MAC ACL by specifying the name of the list. After entering this command, the mode changes to extended MAC ACL configuration.
Step 4	{**permit** | **deny**} {**host** *src-MAC-addr mask* | **any** | **host** *dst-MAC-addr* | *dst-MAC-addr mask*} [*type mask*]	Specify the type of traffic to permit or deny if the conditions are matched, entering the command as many times as necessary. • For *src-MAC-addr*, enter the MAC address of the host from which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the **any** keyword as an abbreviation for *source* 0.0.0, *source-wildcard* 255.255.255, or by using the **host** keyword for *source* 0.0.0. • For *mask*, enter the wildcard bits by placing ones in the bit positions that you want to ignore. • For *dst-MAC-addr*, enter the MAC address of the host to which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the **any** keyword as an abbreviation for *source* 0.0.0, *source-wildcard* 255.255.255, or by using the **host** keyword for *source* 0.0.0. • (Optional) For *type mask*, specify the Ethertype number of a packet with Ethernet II or SNAP encapsulation to identify the protocol of the packet. For *type*, the range is from 0 to 65535, typically specified in hexadecimal. For *mask*, enter the *don't care* bits applied to the Ethertype before testing for a match. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show access-lists** [*access-list-number* | *access-list-name*]	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no mac access-list extended** *access-list-name* global configuration command.

This example shows how to create a Layer 2 MAC ACL with two permit statements. The first statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001. The second statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 to the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-macl)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-macl)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
! (Note: all other access implicitly denied)
```

Classifying Traffic by Using Class Maps

You use the **class-map** global configuration command to isolate a specific traffic flow (or class) from all other traffic and to name it. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can include criterion such as an ACL, IP precedence values, or DSCP values. The match criterion is defined with one match statement entered within the class-map configuration mode.

Note You can also create class-maps during policy map creation by using the **class** policy-map configuration command. For more information, see the "Classifying, Policing, and Marking Traffic by Using Policy Maps" section on page 20-32.

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*] or **access-list** *access-list-number* {**deny** | **permit**} *protocol source* [*source-wildcard*] *destination* [*destination-wildcard*] or **mac access-list extended** *name* {**permit** | **deny**} {**host** *src-MAC-addr mask* | **any** | **host** *dst-MAC-addr* | *dst-MAC-addr mask*} [*type mask*]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the "Classifying Traffic by Using ACLs" section on page 20-27. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.

	Command	Purpose
Step 4	**class-map** *class-map-name* [**match-all** | **match-any**]	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. • For *class-map-name*, specify the name of the class map. • (Optional) Use the **match-all** keyword to perform a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) Use the **match-any** keyword to perform a logical-OR of all matching statements under this class map. One or more match criteria must be matched. If neither the **match-all** or **match-any** keyword is specified, the default is **match-all**. **Note** Because only one **match** command per class map is supported, the **match-all** and **match-any** keywords function the same.
Step 5	**match** {**access-group** *acl-index-or-name* | **ip dscp** *dscp-list* | **ip precedence** *ip-precedence-list*}	Define the match criterion to classify traffic. By default, no match criterion is supported. Only one match criterion per class map is supported, and only one ACL per class map is supported. • For **access-group** *acl-index-or-name*, specify the number or name of the ACL created in Step 3. • For **ip dscp** *dscp-list*, enter a list of up to eight IP DSCP values to match against incoming packets. Separate each value with a space. The range is 0 to 63. • For **ip precedence** *ip-precedence-list*, enter a list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7. **Note** You can use the **match ip dscp** *dscp-list* class-map configuration command only in a policy map that is attached to an egress interface.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show class-map**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an existing class map, use the **no class-map** *class-map-name* [**match-all** | **match-any**] global configuration command. To remove a match criterion, use the **no match** {**access-group** *acl-index-or-name* | **ip dscp** | **ip precedence**} class-map configuration command.

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is access list 103. It permits traffic from any host to any destination that matches a DSCP value of 10.

```
Switch(config)# access-list 103 permit any any dscp 10
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class2*, which matches incoming traffic with DSCP values of 10, 11, and 12.

```
Switch(config)# class-map class2
Switch(config-cmap)# match ip dscp 10 11 12
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class3*, which matches incoming traffic with IP-precedence values of 5, 6, and 7:

```
Switch(config)# class-map class3
Switch(config-cmap)# match ip precedence 5 6 7
Switch(config-cmap)# end
Switch#
```

Classifying, Policing, and Marking Traffic by Using Policy Maps

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.
- A policy-map trust state supersedes an interface trust state.

You can attach only one policy map per interface per direction.

Beginning in privileged EXEC mode, follow these steps to create a policy map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*] or **access-list** *access-list-number* {**deny** | **permit**} *protocol source* [*source-wildcard*] *destination* [*destination-wildcard*] or **mac access-list extended** *access-list name* {**permit** | **deny**} {*source-MAC-addr mask* | **any** | **host**} {*destination-MAC-addr mask* | **any** | **host**} [*ethertype*]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the "Classifying Traffic by Using ACLs" section on page 20-27. **Note** When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.

	Command	Purpose
Step 4	**policy-map** *policy-map-name*	Create a policy map by entering the policy map name, and enter policy-map configuration mode. By default, no policy maps are defined. The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.
Step 5	**class** *class-map-name* [**access-group** *acl-index-or-name* | **dscp** *dscp-list* | **precedence** *ip-precedence-list*]	Define a traffic classification, and enter policy-map class configuration mode. By default, no policy map class-maps are defined. If a traffic class has already been defined by using the **class-map** global configuration command, specify its name for *class-map-name* in this command. To define a class map that uses an access list to filter traffic or that matches traffic to the specified DSCP or IP precedence values, use one of these keywords: • For **access-group** *acl-index-or-name*, specify the number or name of the ACL created in Step 3. • For **dscp** *dscp-list*, specify a list of up to eight IP DSCP values to be matched against incoming packets. Separate each value with a space. The range is 0 to 63. • For **precedence** *ip-precedence-list*, specify a list of up to eight IP-precedence values to be matched against incoming packets. Separate each value with a space. The range is 0 to 7.

	Command	Purpose
Step 6	**trust** [**cos** | **dscp** | **ip-precedence**]	Configure the trust state, which selects the value that QoS uses as the source of the internal DSCP value. **Note** This command is mutually exclusive with the **set** command within the same policy map. If you enter the **trust** command, then skip Step 7. By default, the port is not trusted. If no keyword is specified when the command is entered, the default is **dscp**. The keywords have these meanings: • **cos**—QoS derives the internal DSCP value by using the received or default port CoS value and the CoS-to-DSCP map. • **dscp**—QoS derives the internal DSCP value by using the DSCP value from the ingress packet. For non-IP packets that are tagged, QoS derives the internal DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the internal DSCP value by using the default port CoS value. In either case, the internal DSCP value is derived from the CoS-to-DSCP map. • **ip-precedence**—QoS derives the internal DSCP value by using the IP precedence value from the ingress packet and the IP-precedence-to-DSCP map. For non-IP packets that are tagged, QoS derives the internal DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the internal DSCP value by using the default port CoS value. In either case, the internal DSCP value is derived from the CoS-to-DSCP map. For more information, see the "Configuring the CoS-to-DSCP Map" section on page 20-39.
Step 7	**set** {**ip dscp** *new-dscp* | **ip precedence** *new-precedence*}	Classify IP traffic by setting a new value in the packet. • For **ip dscp** *new-dscp*, enter a new DSCP value to be assigned to the classified traffic. The range is 0 to 63. • For **ip precedence** *new-precedence*, enter a new IP-precedence value to be assigned to the classified traffic. The range is 0 to 7.

	Command	Purpose
Step 8	**police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}]	Define a policer for the classified traffic. You can configure up to 128 policers on ingress Gigabit-capable Ethernet ports, up to 8 policers on ingress 10/100 Ethernet ports, and up to 8 policers on egress ports. • For *rate-bps*, specify average traffic rate in bits per second (bps). The range is 8000 to 2000000000. • For *burst-byte*, specify the normal burst size in bytes. The range is 8000 to 512000000. • (Optional) Specify the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action policed-dscp-transmit** keywords to mark down the DSCP value (by using the policed-DSCP map) and send the packet. For more information, see the "Configuring the Policed-DSCP Map" section on page 20-41.
Step 9	**exit**	Return to policy map configuration mode.
Step 10	**exit**	Return to global configuration mode.
Step 11	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to attach to the policy map. Valid interfaces include physical interfaces.
Step 12	**service-policy** {**input** *policy-map-name* | **output** *policy-map-name*}	Apply a policy map to the input or output of a particular interface. Only one policy map per interface per direction is supported. • Use **input** *policy-map-name* to apply the specified policy-map to the input of an interface. • Use **output** *policy-map-name* to apply the specified policy-map to the output of an interface. **Note** Policy maps that contain **set** or **trust** policy-map class configuration commands or that have ACL classification cannot be attached to an output interface.
Step 13	**end**	Return to privileged EXEC mode.
Step 14	**show policy-map** [*policy-map-name* [**class** *class-name*]]	Verify your entries.
Step 15	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an existing policy map, use the **no policy-map** *policy-map-name* global configuration command. To delete an existing class map, use the **no class** *class-map-name* policy-map configuration command. To return to the default trust state, use the **no trust** [**cos** | **dscp** | **ip-precedence**] policy-map configuration command. To remove an assigned DSCP or IP precedence value, use the **no set** {**ip dscp** *new-dscp* | **ip precedence** *new-precedence*} policy-map configuration command. To remove an existing policer, use the **no police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}] policy-map configuration command. To remove the policy map and interface association, use the **no service-policy** {**input** *policy-map-name* | **output** *policy-map-name*} interface configuration command.

This example shows how to create a policy map and attach it to an ingress interface. In the configuration, the IP standard ACL permits traffic from network 10.1.0.0. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent:

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# policy-map flow1t
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police 48000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input flow1t
```

This example shows how to create a policy map and attach it to an ingress interface. In the configuration, the IP extended ACL permits Transmission Control Protocol (TCP) traffic with an IP precedence of 4 from any host destined for the host at 224.0.0.5. For traffic matching this classification, the DSCP value in the incoming packet is set to 63.

```
Switch(config)# access-list 104 permit tcp any host 224.0.0.5 precedence 4
Switch(config)# policy-map ip104
Switch(config-pmap)# class ipclass104 access-group 104
Switch(config-pmap-c)# set ip dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input ip104
```

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress interface. The first permit statement allows traffic from the host with MAC address 0001.0000.0001 destined for the host with MAC address 0002.0000.0001. The second permit statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 destined for the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit 0001.0000.0003 0.0.0 0002.0000.0003 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0004 0.0.0 0002.0000.0004 0.0.0 aarp
Switch(config-ext-mac)# exit
Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set ip dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set ip dscp 45
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1
```

Classifying, Policing, and Marking Traffic by Using Aggregate Policers

By using an aggregate policer, you can create a policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps or interfaces.

Beginning in privileged EXEC mode, follow these steps to create an aggregate policer:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**mls qos aggregate-police** *aggregate-policer-name rate-bps burst-byte* **exceed-action** {**drop** | **policed-dscp-transmit**}	Define the policer parameters that can be applied to multiple traffic classes within the same policy map. By default, no aggregate policer is defined. You can configure up to 128 policers on ingress Gigabit-capable Ethernet ports, up to 8 policers on ingress 10/100 Ethernet ports, and up to 8 policers on egress ports. • For *aggregate-policer-name*, specify the name of the aggregate policer. • For *rate-bps*, specify average traffic rate in bits per second (bps). The range is 8000 to 2000000000. • For *burst-byte*, specify the normal burst size in bytes. The range is 8000 to 512000000. • (Optional) Specify the action to take when the rates are exceeded. Use the **exceed-action drop** keywords to drop the packet. Use the **exceed-action policed-dscp-transmit** keywords to mark down the DSCP value (by using the policed-DSCP map) and send the packet. For more information, see the "Configuring the Policed-DSCP Map" section on page 20-41.
Step 4	**class-map** *class-map-name* [**match-all** | **match-any**]	Create a class map to classify traffic as necessary. For more information, see the "Classifying Traffic by Using Class Maps" section on page 20-30.
Step 5	**policy-map** *policy-map-name*	Create a policy map by entering the policy map name, and enter policy-map configuration mode. For more information, see the "Classifying, Policing, and Marking Traffic by Using Policy Maps" section on page 20-32.
Step 6	**class** *class-map-name* [**access-group** *acl-index-or-name* | **dscp** *dscp-list* | **precedence** *ip-precedence-list*]	Define a traffic classification, and enter policy-map class configuration mode. For more information, see the "Classifying, Policing, and Marking Traffic by Using Policy Maps" section on page 20-32.
Step 7	**police aggregate** *aggregate-policer-name*	Apply an aggregate policer to multiple classes in the same policy map. For *aggregate-policer-name*, enter the name specified in Step 3.
Step 8	**exit**	Return to global configuration mode.

	Command	Purpose
Step 9	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to attach to the policy map. Valid interfaces include physical interfaces.
Step 10	**service-policy** {**input** *policy-map-name* | **output** *policy-map-name*}	Apply a policy map to the input or output of a particular interface. Only one policy map per interface per direction is supported. • Use **input** *policy-map-name* to apply the specified policy-map to the input of an interface. • Use **output** *policy-map-name* to apply the specified policy-map to the output of an interface. Policy maps that contain **set** or **trust** policy-map class configuration commands or that have ACL classification cannot be attached to an egress interface.
Step 11	**end**	Return to privileged EXEC mode.
Step 12	**show mls qos aggregate-policer** [*aggregate-policer-name*]	Verify your entries.
Step 13	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the specified aggregate policer from a policy map, use the **no police aggregate** *aggregate-policer-name* policy map configuration mode. To delete an aggregate policer and its parameters, use the **no mls qos aggregate-policer** *aggregate-policer-name* global configuration command.

This example shows how to create an aggregate policer and attach it to multiple classes within a policy map. In the configuration, the IP ACLs permit traffic from network 10.1.0.0 and from host 11.3.1.1. For traffic coming from network 10.1.0.0, the DSCP in the incoming packets is trusted. For traffic coming from host 11.3.1.1, the DSCP in the packet is changed to 56. The traffic rate from the 10.1.0.0 network and from host 11.3.1.1 is policed. If the traffic exceeds an average rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent. The policy map is attached to an ingress interface.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# access-list 2 permit 11.3.1.1
Switch(config)# mls qos aggregate-police transmit1 48000 8000 exceed-action
policed-dscp-transmit
Switch(config)# policy-map aggflow1
Switch(config-pmap)# class ipclass1 access-group 1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# class ipclass2 access-group 2
Switch(config-pmap-c)# set ip dscp 56
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input aggflow1
Switch(config-if)# exit
```

Configuring DSCP Maps

This section describes how to configure the DSCP maps. It contains this configuration information:

- Configuring the CoS-to-DSCP Map, page 20-39
- Configuring the IP-Precedence-to-DSCP Map, page 20-40
- Configuring the Policed-DSCP Map, page 20-41
- Configuring the DSCP-to-CoS Map, page 20-42
- Configuring the DSCP-to-DSCP-Mutation Map, page 20-43

All the maps, except the DSCP-to-DSCP-mutation map, are globally defined and are applied to all ports.You can have multiple DSCP-to-DSCP-mutation maps and apply them to different Gigabit-capable Ethernet ports. However, on 10/100 Ethernet ports, you can attach only one DSCP-to-DSCP-mutation map to a group of twelve ports.

Configuring the CoS-to-DSCP Map

You use the CoS-to-DSCP map to map CoS values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 20-3 shows the default CoS-to-DSCP map.

Table 20-3 Default CoS-to-DSCP Map

CoS value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the CoS-to-DSCP map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map cos-dscp** *dscp1...dscp8*	Modify the CoS-to-DSCP map. For *dscp1...dscp8*, enter 8 DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps cos-dscp**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos cos-dscp** global configuration command.

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch# configure terminal
Switch(config)# mls qos map cos-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps cos-dscp

Cos-dscp map:
        cos:   0  1  2  3  4  5  6  7
     --------------------------------
       dscp:   10 15 20 25 30 35 40 45
```

Configuring the IP-Precedence-to-DSCP Map

You use the IP-precedence-to-DSCP map to map IP precedence values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 20-4 shows the default IP-precedence-to-DSCP map:

Table 20-4 Default IP-Precedence-to-DSCP Map

IP precedence value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the IP-precedence-to-DSCP map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map ip-prec-dscp** *dscp1...dscp8*	Modify the IP-precedence-to-DSCP map. For *dscp1...dscp8*, enter 8 DSCP values that correspond to the IP precedence values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps ip-prec-dscp**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos ip-prec-dscp** global configuration command.

This example shows how to modify and display the IP-precedence-to-DSCP map:

```
Switch# configure terminal
Switch(config)# mls qos map ip-prec-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps ip-prec-dscp

IpPrecedence-dscp map:
     ipprec:   0  1  2  3  4  5  6  7
     --------------------------------
       dscp:   10 15 20 25 30 35 40 45
```

Configuring the Policed-DSCP Map

You use the policed-DSCP map to mark down a DSCP value to a new value as the result of a policing and marking action.

The default policed-DSCP map is a null map, which maps an incoming DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the policed-DSCP map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map policed-dscp** *dscp-list* **to** *mark-down-dscp*	Modify the policed-DSCP map. • For *dscp-list*, enter up to 8 DSCP values separated by spaces. Then enter the **to** keyword. • For *mark-down-dscp*, enter the corresponding policed (marked down) DSCP value.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps policed-dscp**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos policed-dscp** global configuration command.

This example shows how to map DSCP 50 to 57 to a marked-down DSCP value of 0:

```
Switch# configure terminal
Switch(config)# mls qos map policed-dscp 50 51 52 53 54 55 56 57 to 0
Switch(config)# end
Switch# show mls qos maps policed-dscp
Policed-dscp map:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 01 02 03 04 05 06 07 08 09
      1 :    10 11 12 13 14 15 16 17 18 19
      2 :    20 21 22 23 24 25 26 27 28 29
      3 :    30 31 32 33 34 35 36 37 38 39
      4 :    40 41 42 43 44 45 46 47 48 49
      5 :    00 00 00 00 00 00 00 00 58 59
      6 :    60 61 62 63
```

Note In this policed-DSCP map, the marked-down DSCP values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the marked-down value. For example, an original DSCP value of 53 corresponds to a marked-down DSCP value of 0.

Configuring the DSCP-to-CoS Map

You use the DSCP-to-CoS map to generate a CoS value, which is used to select one of the four egress queues.

Table 20-5 shows the default DSCP-to-CoS map.

Table 20-5 Default DSCP-to-CoS Map

DSCP value	0–7	8–15	16–23	24–31	32–39	40–47	48–55	56–63
CoS value	0	1	2	3	4	5	6	7

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-CoS map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map dscp-cos** *dscp-list* **to** *cos*	Modify the DSCP-to-CoS map. • For *dscp-list*, enter up to 8 DSCP values separated by spaces. Then enter the **to** keyword. • For *cos*, enter the CoS value to which the DSCP values correspond. The DSCP range is 0 to 63; the CoS range is 0 to 7.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show mls qos maps dscp-to-cos**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos dscp-cos** global configuration command.

This example shows how to map DSCP values 0, 8, 16, 24, 32, 40, 48, and 50 to CoS value 0 and to display the map:

```
Switch# configure terminal
Switch(config)# mls qos map dscp-cos 0 8 16 24 32 40 48 50 to 0
Switch(config)# end
Switch# show mls qos maps dscp-cos
Dscp-cos map:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 00 00 00 00 00 00 00 00 01
      1 :    01 01 01 01 01 01 00 02 02 02
      2 :    02 02 02 02 00 03 03 03 03 03
      3 :    03 03 00 04 04 04 04 04 04 04
      4 :    00 05 05 05 05 05 05 05 00 06
      5 :    00 06 06 06 06 06 07 07 07 07
      6 :    07 07 07 07
```

Note In the above DSCP-to-CoS map, the CoS values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the DSCP; the d2 row specifies the least-significant digit of the DSCP. The intersection of the d1 and d2 values provides the CoS value. For example, in the DSCP-to-CoS map, a DSCP value of 08 corresponds to a CoS value of 0.

Configuring the DSCP-to-DSCP-Mutation Map

You apply the DSCP-to-DSCP-mutation map to a port at the boundary of a QoS administrative domain. If the two domains have different DSCP definitions between them, you use the DSCP-to-DSCP-mutation map to translate a set of DSCP values to match the definition of the other domain.

The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-DSCP-mutation map:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos map dscp-mutation** *dscp-mutation-name in-dscp* **to** *out-dscp*	Modify the DSCP-to-DSCP-mutation map. • For *dscp-mutation-name*, enter the mutation map name. You can create more than one map by specifying a new name. • For *in-dscp*, enter up to 8 DSCP values separated by spaces. Then enter the **to** keyword. • For *out-dscp*, enter up to 8 DSCP values separated by spaces. The DSCP range is 0 to 63.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to which to attach the map. Valid interfaces include physical interfaces.
Step 4	**mls qos trust dscp**	Configure the ingress port as a DSCP-trusted port.
Step 5	**mls qos dscp-mutation** *dscp-mutation-name*	Apply the map to the specified ingress DSCP-trusted port. For *dscp-mutation-name*, enter the mutation map name specified in Step 2. You can apply the map to different Gigabit-capable Ethernet ports. However, on 10/100 Ethernet ports, you can attach only one DSCP-to-DSCP-mutation map to a group of twelve ports. For example, Fast Ethernet ports 0/1 to 0/12 are a group, Fast Ethernet ports 0/13 to 0/24 are a group, Gigabit Ethernet port 0/1 is a group, and Gigabit Ethernet port 0/2 is a group. When applying a mutation map to any port in a group, all ports in the same group are automatically configured with the same map.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show mls qos maps dscp-mutation**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos dscp-mutation** *dscp-mutation-name* global configuration command.

This example shows how to define the DSCP-to-DSCP-mutation map. All the entries that are not explicitly configured are not modified (remains as specified in the null map):

```
Switch# configure terminal
Switch(config)# mls qos map dscp-mutation mutation1 1 2 3 4 5 6 7 to 0
Switch(config)# mls qos map dscp-mutation mutation1 8 9 10 11 12 13 to 10
Switch(config)# mls qos map dscp-mutation mutation1 20 21 22 to 20
Switch(config)# mls qos map dscp-mutation mutation1 30 31 32 33 34 to 30
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation mutation1
Switch(config-if)# end
Switch# show mls qos maps dscp-mutation mutation1
Dscp-dscp mutation map:
   mutation1:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 00 00 00 00 00 00 00 10 10
      1 :    10 10 10 10 14 15 16 17 18 19
      2 :    20 20 20 23 24 25 26 27 28 29
      3 :    30 30 30 30 30 35 36 37 38 39
      4 :    40 41 42 43 44 45 46 47 48 49
      5 :    50 51 52 53 54 55 56 57 58 59
      6 :    60 61 62 63
```

Note In the above DSCP-to-DSCP-mutation map, the mutated values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the mutated value. For example, a DSCP value of 12 corresponds to a mutated value of 10.

Configuring Egress Queues on Gigabit-Capable Ethernet Ports

This section describes how to configure the egress queues on Gigabit-capable Ethernet ports. For information on configuring 10/100 Ethernet ports, see "Configuring Egress Queues on 10/100 Ethernet Ports" section on page 20-51.

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are assigned (by CoS value) to each queue?
- How much of the available buffer space (limit) is allotted to each queue?
- What drop percentage thresholds apply to each queue and which DSCP values map to each threshold?
- Is one of the queues the expedite (high-priority) egress queue?
- How much of the available bandwidth is allotted to each queue?

This section contains this configuration information:

- Mapping CoS Values to Select Egress Queues, page 20-45
- Configuring the Egress Queue Size Ratios, page 20-46
- Configuring Tail-Drop Threshold Percentages, page 20-47
- Configuring WRED Drop Thresholds Percentages, page 20-48
- Configuring the Egress Expedite Queue, page 20-50
- Allocating Bandwidth among Egress Queues, page 20-50

Mapping CoS Values to Select Egress Queues

Beginning in privileged EXEC mode, follow these steps to map CoS ingress values to select one of the egress queues:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.
Step 4	**wrr-queue cos-map** *queue-id cos1 ... cos8*	Map assigned CoS values to select one of the egress queues. The default map has these values: CoS value 0, 1 selects queue 1. CoS value 2, 3 selects queue 2. CoS value 4, 5 selects queue 3. CoS value 6, 7 selects queue 4. • For *queue-id*, specify the ID of the egress queue. The range is 1 to 4, where 4 can be configured as the expedite queue. For more information, see the "Configuring the Egress Expedite Queue" section on page 20-50. • For *cos1 ... cos8*, specify the CoS values that select a queue. Enter up to eight CoS values. Separate each value with a space. The range is 0 to 7.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface queueing**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the default CoS-to-egress-queue map, use the **no wrr-queue cos-map** [*queue-id*] interface configuration command.

This example shows how to map CoS values 6 and 7 to queue 1, 4 and 5 to queue 2, 2 and 3 to queue 3, 0 and 1 to queue 4.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# wrr-queue cos-map 1 6 7
Switch(config-if)# wrr-queue cos-map 2 4 5
Switch(config-if)# wrr-queue cos-map 3 2 3
Switch(config-if)# wrr-queue cos-map 4 0 1
```

Configuring the Egress Queue Size Ratios

Beginning in privileged EXEC mode, follow these steps to configure the egress queue size ratios:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.
Step 4	**wrr-queue queue-limit** *weight1 weight2 weight3 weight4*	Configure the egress queue size ratios. The defaults weights are 25 (1/4 of the buffer size is allocated to each queue). For *weight1*, *weight2*, *weight3*, and *weight4*, specify a weight from 1 to 100. Separate each value with a space. The relative size difference in the numbers indicates the relative differences in the queue sizes. The port enters a *halt* mode when this command is issued.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface buffers**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default weights, use the **no wrr-queue queue-limit** interface configuration command.

This example shows how to configure the size ratio of the four queues. The ratio of the size allocated for each queue is 1/10, 2/10, 3/10, and 4/10 to queue 1, 2, 3, and 4. (Queue 4 is four times larger than queue 1, twice as large as queue 2, and 1.33 times as large as queue 3.)

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# wrr-queue queue-limit 1 2 3 4
```

Configuring Tail-Drop Threshold Percentages

Tail drop is the default congestion-avoidance technique on Gigabit-capable Ethernet ports. With tail drop, packets are queued until the thresholds are exceeded. For example, all packets with DSCPs assigned to the first threshold are dropped until the threshold is no longer exceeded. However, packets assigned to a second threshold continue to be queued and sent as long as the second threshold is not exceeded.

You modify the DSCP-to-threshold map to determine which DSCPs are mapped to which threshold ID by using the **wrr-queue dscp-map** interface configuration command. By default, all DSCPs are mapped to threshold 1, and when this threshold is exceeded, all the packets are dropped.

If you use tail-drop thresholds, you cannot use WRED, and vice versa.

Beginning in privileged EXEC mode, follow these steps to configure the tail-drop threshold percentage values on Gigabit-capable Ethernet ports:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.
Step 4	**wrr-queue threshold** *queue-id threshold-percentage1 threshold-percentage2*	Configure tail-drop threshold percentages on each egress queue. The default threshold is 100 percent for thresholds 1 and 2. • For *queue-id*, specify the ID of the egress queue. The range is 1 to 4. • For *threshold-percentage1 threshold-percentage2*, specify the tail-drop threshold percentage values. Separate each value with a space. The range is 1 to 100.
Step 5	**exit**	Return to global configuration mode.
Step 6	**interface** *interface-id*	Enter interface configuration mode, and specify the ingress Gigabit-capable Ethernet interface.
Step 7	**wrr-queue dscp-map** *threshold-id dscp1 ... dscp8*	Map DSCP values to the tail-drop thresholds of the egress queues. By default, all DSCP values are mapped to threshold 1. • For *threshold-id*, specify the threshold ID of the queue. The range is 1 to 2. • For *dscp1 ... dscp8*, specify the DSCP values that are mapped to the threshold ID. Enter up to eight DSCP values per command. Separate each value with a space. The range is 0 to 63.
Step 8	**end**	Return to privileged EXEC mode.
Step 9	**show running-config** or **show mls qos interface** *interface-id* **queueing**	Verify the DSCP-to-threshold map.
Step 10	**show mls qos interface buffers**	Verify the thresholds.
Step 11	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default thresholds, use the **no wrr-queue threshold** *queue-id* interface configuration command. To return to the default DSCP-to-threshold map, use the **no wrr-queue dscp-map** [*threshold-id*] interface configuration command.

This example shows how to configure the tail-drop queue threshold values for queue 1 to 10 percent and 100 percent, for queue 2 to 40 percent and 100 percent, for queue 3 to 60 percent and 100 percent, and for queue 4 to 80 percent and 100 percent on the egress interface (Gigabit Ethernet 0/1). The ingress interface (Gigabit Ethernet 0/2) is configured to trust the DSCP in the incoming packets, to map DSCPs 0, 8, 16, 24, 32, 40, 48, and 56 to threshold 1, and to map DSCPs 10, 20, 30, 40, 50, and 60 to threshold 2:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# wrr-queue threshold 1 10 100
Switch(config-if)# wrr-queue threshold 2 40 100
Switch(config-if)# wrr-queue threshold 3 60 100
Switch(config-if)# wrr-queue threshold 4 80 100
Switch(config-if)# exit
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# mls qos trust dscp
Switch(config-if)# wrr-queue dscp-map 1 0 8 16 24 32 40 48 56
Switch(config-if)# wrr-queue dscp-map 2 10 20 30 40 50 60
```

As a result of this configuration, when queue 1 is filled above 10 percent, packets with DSCPs 0, 8, 16, 24, 32, 40, 48, and 56 are dropped. The same packets are dropped when queue 2 is filled above 40 percent, queue 3 above 60 percent, and queue 4 above 80 percent. When the second threshold (100 percent) is exceeded, all queues drop packets with DSCPs 10, 20, 30, 40, 50, and 60.

Configuring WRED Drop Thresholds Percentages

WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once.

All packets with DSCPs assigned to the first threshold are randomly dropped when the first threshold is exceeded. However, packets with DSCPs assigned to the second threshold continue to be queued and sent as long as the second threshold is not exceeded. Each threshold percentage represents where WRED starts to randomly drop packets. By default, WRED is disabled.

Beginning in privileged EXEC mode, follow these steps to configure the WRED drop threshold percentage values on Gigabit-capable Ethernet ports:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.

	Command	Purpose
Step 4	**wrr-queue random-detect max-threshold** *queue-id threshold-percentage1 threshold-percentage2*	Configure WRED drop threshold percentages on each egress queue. The default, WRED is disabled, and no thresholds are configured. • For *queue-id*, specify the ID of the egress queue. The range is 1 to 4, where queue 4 can be configured as the expedite queue. For more information, see the "Configuring the Egress Expedite Queue" section on page 20-50. • For *threshold-percentage1 threshold-percentage2*, specify the threshold percentage values. Separate each value with a space. The range is 1 to 100.
Step 5	**exit**	Return to global configuration mode.
Step 6	**interface** *interface-id*	Enter interface configuration mode, and specify the ingress Gigabit-capable Ethernet interface.
Step 7	**wrr-queue dscp-map** *threshold-id dscp1 ... dscp8*	Map DSCP values to the WRED drop thresholds of the egress queues. By default, all DSCP values are mapped to threshold 1. • For *threshold-id*, specify the threshold ID of the queue. The range is 1 to 2. • For *dscp1 ... dscp8*, specify the DSCP values that are mapped to the threshold ID. Enter up to eight DSCP values per command. Separate each value with a space. The range is 0 to 63.
Step 8	**show running-config**	Verify the DSCP-to-threshold map.
Step 9	**show mls qos interface buffers**	Verify the thresholds.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable WRED, use the **no wrr random-detect max-threshold** [*queue-id*] interface configuration command. To return to the default DSCP-to-threshold map, use the **no wrr-queue dscp-map** [*threshold-id*] interface configuration command.

This example shows how to configure the WRED queue threshold values for queue 1 to 50 percent and 100 percent, for queue 2 to 70 percent and 100 percent, for queue 3 to 50 percent and 100 percent, and for queue 4 to 70 percent and 100 percent on the egress interface (Gigabit Ethernet 0/1). The ingress interface (Gigabit Ethernet 0/2) is configured to trust the DSCP in the incoming packets, to map DSCPs 0, 8, 16, 24, 32, 40, 48, and 56 to threshold 1, and to map DSCPs 10, 20, 30, 40, 50, and 60 to threshold 2.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# wrr-queue random-detect max-threshold 1 50 100
Switch(config-if)# wrr-queue random-detect max-threshold 2 70 100
Switch(config-if)# wrr-queue random-detect max-threshold 3 50 100
Switch(config-if)# wrr-queue random-detect max-threshold 4 70 100
Switch(config-if)# exit
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# mls qos trust dscp
Switch(config-if)# wrr-queue dscp-map 1 0 8 16 24 32 40 48 56
Switch(config-if)# wrr-queue dscp-map 2 10 20 30 40 50 60
```

As a result of this configuration, when the queues 1 and 3 are filled above 50 percent, packets with DSCPs 0, 8, 16, 24, 32, 40, 48, and 56 are randomly dropped. The same packets are randomly dropped when queues 2 and 4 are filled above 70 percent. When the second threshold (100 percent) is exceeded, all queues randomly drop packets with DSCPs 10, 20, 30, 40, 50, and 60.

Configuring the Egress Expedite Queue

You can ensure that certain packets have priority over all others by queuing them in the egress expedite queue. This queue is serviced until it is empty and before the other queues are serviced.

Beginning in privileged EXEC mode, follow these steps to enable the egress expedite queue:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.
Step 4	**priority-queue out**	Enable the egress expedite queue, which is disabled by default. When you configure this command, the WRR weight and queue size ratios are affected because there is one fewer queue participating in WRR. This means that *weight4* in the **wrr-queue bandwidth** command is ignored (not used in the ratio calculation).
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the egress expedite queue, use the **no priority-queue out** interface configuration command.

Allocating Bandwidth among Egress Queues

You need to specify how much of the available bandwidth is allocated to each queue. The ratio of the weights is the ratio of frequency in which the WRR scheduler drops packets from each queue.

Beginning in privileged EXEC mode, follow these steps to allocate bandwidth to each queue:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress Gigabit-capable Ethernet interface.

	Command	Purpose
Step 4	**wrr-queue bandwidth** *weight1 weight2 weight3 weight4*	Assign WRR weights to the egress queues. By default, all the weights are set to 25 (1/4 of the bandwidth is allocated to each queue). For *weight1 weight2 weight3 weight4*, enter the ratio, which determines the ratio of the frequency in which the WRR scheduler drops packets. Separate each value with a space. The range is 0 to 65536. All four queues participate in the WRR unless the expedite queue (queue 4) is enabled, in which case *weight4* is ignored (not used in the ratio calculation). The expedite queue is a strict-priority queue, and it is serviced until empty before the other queues are serviced. A weight of 0 means no bandwidth is allocated for that queue, and the available bandwidth is shared among the remaining queues.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface queueing**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default bandwidth setting, use the **no wrr-queue bandwidth** interface configuration command.

This example shows how to configure the weight ratio of the WRR scheduler running on the egress queues. In this example, four queues are used (no expedite queue), and the ratio of the bandwidth allocated for each queue is 1/(1+2+3+4), 2/(1+2+3+4), 3/(1+2+3+4), and 4/(1+2+3+4), which is 1/10, 1/5, 3/10, and 2/5 for queues 1, 2, 3, and 4.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# wrr-queue bandwidth 1 2 3 4
```

Configuring Egress Queues on 10/100 Ethernet Ports

This section describes how to configure the egress queues on 10/100 Ethernet ports. For information on configuring Gigabit-capable Ethernet ports, see the "Configuring Egress Queues on Gigabit-Capable Ethernet Ports" section on page 20-44.

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are assigned (by CoS value) to each queue?
- How much of the available buffer space is allotted to each queue?
- Is one of the queues the expedite (high-priority) egress queue?
- How much of the available bandwidth is allotted to each queue?

This section contains this configuration information:

- Mapping CoS Values to Select Egress Queues, page 20-52
- Configuring the Minimum-Reserve Levels, page 20-53
- Configuring the Egress Expedite Queue, page 20-54
- Allocating Bandwidth among Egress Queues, page 20-54

Mapping CoS Values to Select Egress Queues

Beginning in privileged EXEC mode, follow these steps to map CoS ingress values to select one of the egress queues:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress 10/100 Ethernet interface.
Step 4	**wrr-queue cos-map** *queue-id cos1 ... cos8*	Map assigned CoS values to select one of the egress queues. Theses are the default map values: CoS value 0, 1 selects queue 1. CoS value 2, 3 selects queue 2. CoS value 4, 5 selects queue 3. CoS value 6, 7 selects queue 4. • For *queue-id*, specify the ID of the egress queue. The range is 1 to 4, where 4 can be configured as the expedite queue. For more information, see the "Configuring the Egress Expedite Queue" section on page 20-54. • For *cos1 ... cos8*, specify the CoS values that select a queue. Enter up to eight CoS values. Separate each value with a space. The range is 0 to 7.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface queueing**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default CoS-to-egress-queue map, use the **no wrr-queue cos-map** [*queue-id*] interface configuration command.

This example shows how to map CoS values 6 and 7 to queue 1, 4 and 5 to queue 2, 2 and 3 to queue 3, and 0 and 1 to queue 4.

```
Switch(config)# interface fastethernet0/1
Switch(config-if)# wrr-queue cos-map 1 6 7
Switch(config-if)# wrr-queue cos-map 2 4 5
Switch(config-if)# wrr-queue cos-map 3 2 3
Switch(config-if)# wrr-queue cos-map 4 0 1
```

Configuring the Minimum-Reserve Levels

You can configure the buffer size of the minimum-reserve levels on all 10/100 ports and assign the minimum-reserve level to an egress queue on a 10/100 Ethernet port.

Beginning in privileged EXEC mode, follow these steps to configure the egress queue sizes:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**mls qos min-reserve** *min-reserve-level min-reserve-buffersize*	Configure the buffer size of the minimum-reserve level, if necessary, for all the 10/100 Ethernet ports. By default, the buffer size for all eight minimum-reserve levels is set to 100 packets. • For *min-reserve-level*, specify the minimum-reserve level number. The range is 1 to 8. • For *min-reserve-buffersize*, specify the buffer size. The range is 10 to 170 packets.
Step 4	**interface** *interface-id*	Enter interface configuration mode, and specify the egress 10/100 Ethernet interface.
Step 5	**wrr-queue min-reserve** *queue-id min-reserve-level*	Assign a minimum-reserve level number to a particular egress queue. By default, queue 1 selects minimum-reserve level 1, queue 2 selects minimum-reserve level 2, queue 3 selects minimum-reserve level 3, and queue 4 selects minimum-reserve level 4. • For *queue-id*, specify the ID of the egress queue. The range is 1 to 4, where 4 can be configured as the expedite queue. For more information, see the "Configuring the Egress Expedite Queue" section on page 20-54. • For *min-reserve-level*, specify the minimum-reserve level configured in Step 3.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show mls qos interface buffers**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default minimum-reserve buffer size, use the **no mls qos min-reserve** *min-reserve-level* global configuration command. To return to the default queue selection of the minimum-reserve level, use the **no wrr-queue min-reserve** *queue-id* interface configuration command.

This example shows how to configure minimum-reserve level 5 to 20 packets and to assign minimum-reserve level 5 to egress queue 1 on the Fast Ethernet 0/1 interface 0/1:

```
Switch(config)# mls qos min-reserve 5 20
Switch(config)# interface fastethernet0/1
Switch(config-if)# wrr-queue min-reserve 1 5
```

Configuring the Egress Expedite Queue

You can ensure that certain packets have priority over all others by queuing them in the egress expedite queue. This queue is serviced until it is empty and before the other queues are serviced.

Beginning in privileged EXEC mode, follow these steps to enable the egress expedite queue:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress 10/100 Ethernet interface.
Step 4	**priority-queue out**	Enable the egress expedite queue, which is disabled by default. When you configure this command, the WRR weight is affected because there is one fewer queue participating in WRR. This means that *weight4* in the **wrr-queue bandwidth** command is ignored (not used in the ratio calculation).
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the egress expedite queue, use the **no priority-queue out** interface configuration command.

Allocating Bandwidth among Egress Queues

You need to specify how much of the available bandwidth is allocated to each queue. The ratio of the weights is the ratio of frequency in which the WRR scheduler drops packets from each queue.

Beginning in privileged EXEC mode, follow these steps to allocate bandwidth to each queue:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the egress 10/100 Ethernet interface.

	Command	Purpose
Step 4	**wrr-queue bandwidth** *weight1 weight2 weight3 weight4*	Assign WRR weights to the egress queues. By default, all the weights are set to 25 (1/4 of the bandwidth is allocated to each queue). For *weight1 weight2 weight3 weight4*, enter the ratio, which determines the ratio of the frequency in which the WRR scheduler drops packets. Separate each value with a space. The range is 0 to 65536. All four queues participate in the WRR unless the expedite queue (queue 4) is enabled, in which case *weight4* is ignored (not used in the ratio calculation). The expedite queue is a strict-priority queue, and it is serviced until empty before the other queues are serviced. A weight of 0 means no bandwidth is allocated for that queue, and the available bandwidth is shared among the remaining queues.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show mls qos interface queueing**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default bandwidth setting, use the **no wrr-queue bandwidth** interface configuration command.

This example shows how to configure the weight ratio of the WRR scheduler running on the egress queues. In this example, four queues are used (no expedite queue), and the ratio of the bandwidth allocated for each queue is 1/(1+2+3+4), 2/(1+2+3+4), 3/(1+2+3+4), and 4/(1+2+3+4), which is 1/10, 2/10, 3/10, and 4/10 for queues 1, 2, 3, and 4.

```
Switch(config)# interface fastethernet0/1
Switch(config-if)# wrr-queue bandwidth 1 2 3 4
```

Displaying QoS Information

To display the current QoS information, use one or more of the privileged EXEC commands in Table 20-6:

Table 20-6 Commands for Displaying QoS Information

Command	Purpose
show class-map [*class-map-name*]	Display QoS class maps, which define the match criteria to classify traffic.
show policy-map [*policy-map-name* [**class** *class-name*] | [**interface** *interface-id*]	Display QoS policy maps, which define classification criteria for incoming traffic.
show mls qos aggregate-policer [*aggregate-policer-name*]	Display the aggregate policer configuration.
show mls qos maps [**cos-dscp** | **dscp-cos** | **dscp-mutation** | **ip-prec-dscp** | **policed-dscp**]	Display QoS mapping information. Maps are used to generate an internal DSCP value, which represents the priority of the traffic.
show mls qos interface [*interface-id*] [**buffers** | **policers** | **queueing** | **statistics**]	Display QoS information at the interface level, including the configuration of the egress queues and the CoS-to-egress-queue map, which interfaces have configured policers, and ingress and egress statistics (including the number of bytes dropped).[1]

1. You can define up to 16 DSCP values for which byte or packet statistics are gathered by hardware by using the **mls qos monitor** {**bytes** | **dscp** *dscp1 ... dscp8* | **packets**} interface configuration command and the **show mls qos interface statistics** privileged EXEC command.

QoS Configuration Examples

This section provides a QoS migration path to help you quickly implement QoS features based on your existing network and planned changes to your network, as shown in Figure 20-10. It contains this information:

- QoS Configuration for the Common Wiring Closet, page 20-57
- QoS Configuration for the Intelligent Wiring Closet, page 20-58
- QoS Configuration for the Distribution Layer, page 20-59

Figure 20-10 QoS Configuration Example Network

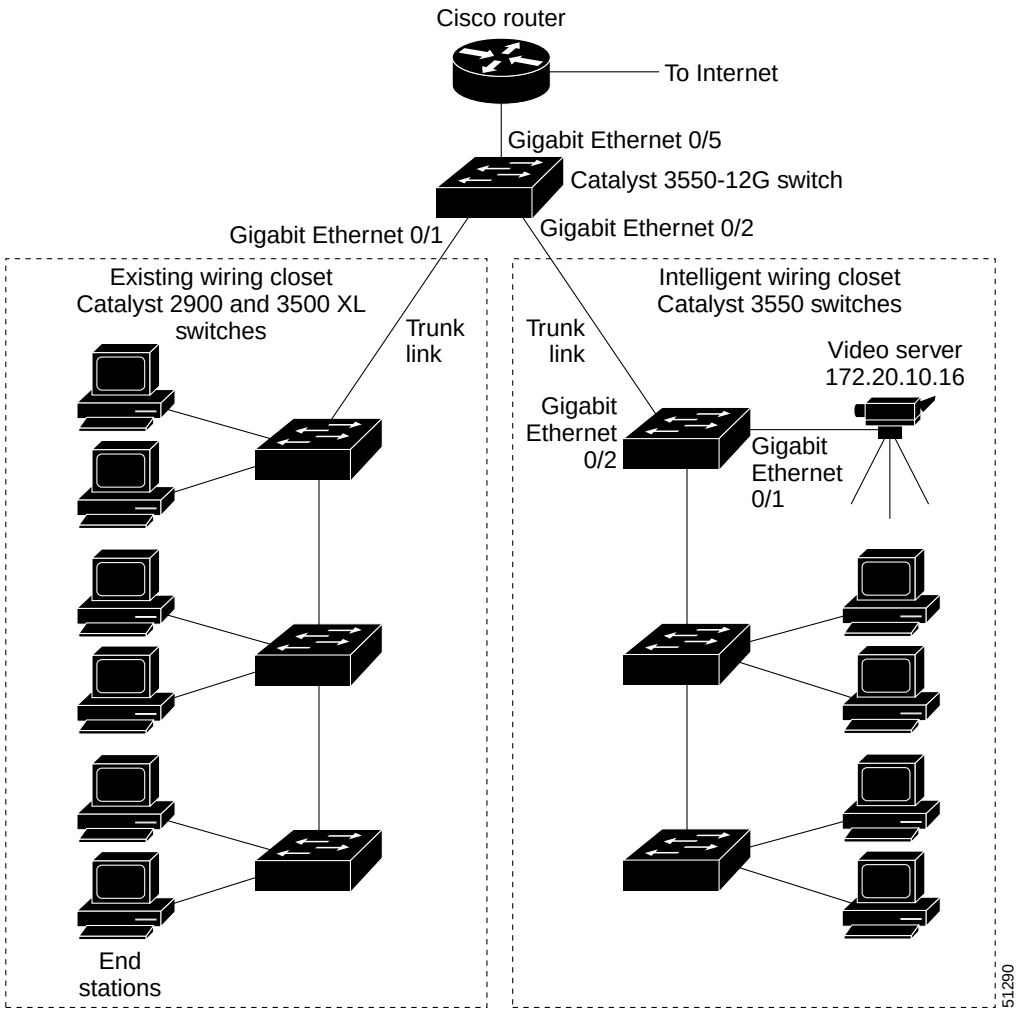

QoS Configuration for the Common Wiring Closet

The common wiring closet in Figure 20-10 consists of existing Catalyst 3500 XL and 2900 XL switches. These switches are running IOS release 12.0(5)XP or later, which supports the QoS-based IEEE 802.1P CoS values. QoS classifies frames by assigning priority-indexed CoS values to them and gives preference to higher-priority traffic.

Recall that on the Catalyst 3500 XL and 2900 XL switches, you can classify untagged (native) Ethernet frames at the ingress ports by setting a default CoS priority (**switchport priority default** *default-priority-id* interface configuration command) for each port. For ISL or IEEE 802.1Q frames with tag information, the priority value from the header frame is used. On the Catalyst 3524-PWR XL and 3548 XL switches, you can override this priority with the default value by using the **switchport priority default override** interface configuration command. For Catalyst 3500 XL, 2950, other 2900 XL models that do not have the override feature, the Catalyst 3550-12T switch at the distribution layer can override the 802.1P CoS value by using the **mls qos cos override** interface configuration command.

For the Catalyst 3500 XL and 2900 XL switches, CoS configures each egress port with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority queue are forwarded only after frames in the high-priority queue are forwarded. Frames that have 802.1P CoS values of 0 to 3 are placed in the normal-priority transmit queue whereas frames with CoS values of 4 to 7 are placed in the expedite (high-priority) queue.

QoS Configuration for the Intelligent Wiring Closet

The intelligent wiring closet in Figure 20-10 is composed of Catalyst 3550 multilayer switches. One of the switches is connected to a video server, which has an IP address of 172.20.10.16.

The object of this example is to prioritize the video traffic over all other traffic. To do so, a DSCP of 56 is assigned to the video traffic. This traffic is stored in the expedite queue (queue 4), which is serviced until empty before the other queues are serviced. The appropriate CoS value selects queue 4 in the CoS-to-egress-queue map.

Beginning in privileged EXEC mode, follow these steps to configure the switch to prioritize video packets over all other traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list 1 permit 172.20.10.16**	Define an IP standard ACL, and permit traffic from the video server at 172.20.10.16.
Step 3	**class-map videoclass**	Create a class map called *videoclass*, and enter class-map configuration mode.
Step 4	**match access-group 1**	Define the match criterion by matching the traffic specified by access list 1.
Step 5	**exit**	Return to global configuration mode.
Step 6	**policy-map videopolicy**	Create a policy map called *videopolicy*, and enter policy-map configuration mode.
Step 7	**class videoclass**	Specify the class on which to act, and enter policy-map class configuration mode.
Step 8	**set ip dscp 56**	For traffic matching ACL 1, set the DSCP of incoming packets to 56.
Step 9	**police 5000000 2000000 exceed-action drop**	Define a policer for the classified video traffic to drop traffic that exceeds 5-Mbps average traffic rate with a 2-MB burst size.
Step 10	**exit**	Return to policy-map configuration mode.
Step 11	**exit**	Return to global configuration mode.
Step 12	**interface gigabitethernet0/1**	Enter interface configuration mode, and specify the ingress interface.
Step 13	**service-policy input videopolicy**	Apply the policy to the ingress interface.
Step 14	**exit**	Return to global configuration mode.
Step 15	**interface gigabitethernet0/2**	Enter interface configuration mode, and specify the egress interface (to configure the queues).
Step 16	**priority-queue out**	Enable the expedite queue.

	Command	Purpose
Step 17	**wrr-queue cos-map 4 6 7**	Configure the CoS-to-egress-queue map so that CoS values 6 and 7 select queue 4 (this is the default setting). Because the default DSCP-to-CoS map has DSCP values 56 to 63 mapped to CoS value 7, the matched traffic that is set to DSCP 56 goes to the queue 4, the priority queue.
Step 18	**end**	Return to privileged EXEC mode.
Step 19	**show class-map videoclass** **show policy-map videopolicy** **show mls qos maps** [**cos-dscp** | **dscp-cos**] **show mls qos interface queueing**	Verify your entries.
Step 20	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

QoS Configuration for the Distribution Layer

This example focuses on the configuration steps for the Catalyst 3550-12G multilayer switch at the distribution layer (see Figure 20-10). Because the classification was performed by the switches at the edge of the network, fewer classification steps are needed at the distribution layer switch.

For the connection to the common wiring closet, Gigabit Ethernet interface 0/1 on the multilayer switch is configured to trust the received CoS value. In this situation, the default CoS-to-DSCP map on the multilayer switch is sufficient. For information on the default map settings, see the "Configuring the CoS-to-DSCP Map" section on page 20-39.

For the connection to the intelligent wiring closet, Gigabit Ethernet interface 0/2 on the multilayer switch is configured to trust the received DSCP value. The DSCP-to-threshold map also needs to be configured on this ingress interface so that on the egress interface, WRED can provide congestion avoidance control. By default, all DSCP values are mapped to threshold 1.

You need to configure several of the switch maps from their default settings. The object of the configuration is to have only DSCP value 56 sent to the expedite queue (queue 4). The default CoS-to-egress-queue map is sufficient; however, you need to configure the DSCP-to-CoS map so that DSCP values 57 to 63 map to CoS 5.

For the egress interface, Gigabit Ethernet interface 0/5, WRR weights need to be configured by using the **wrr-queue bandwidth** interface configuration command. WRED needs to be enabled and the threshold percentages configured for each queue. The bandwidth allocated to each queue must be configured to determine the ratio of the frequency at which packets are dropped from the queue.

Beginning in privileged EXEC mode, follow these steps to configure the multilayer switch at the distribution layer:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**mls qos**	Enable QoS on the switch.
Step 3	**interface gigabitethernet0/1**	Enter interface configuration mode, and specify the ingress interface that is connected to the common wiring closet.
Step 4	**mls qos trust cos**	Classify incoming packets on this port by using the packet CoS value.

	Command	Purpose
Step 5	**switchport mode trunk**	Configure this port as a trunk port.
Step 6	**exit**	Return to global configuration mode.
Step 7	**interface gigabitethernet0/2**	Enter interface configuration mode, and specify the ingress interface connected to the intelligent wiring closet.
Step 8	**mls qos trust dscp**	Classify incoming packets on this port by using the packet DSCP value.
Step 9	**wrr-queue dscp-map** *threshold-id dscp1 ... dscp8*	Map the ingress DSCP values to the WRED thresholds of the egress queues. In the default DSCP-to-threshold map, all DSCP values are mapped to threshold 1. • For *threshold-id*, specify the threshold ID of the queue. The range is 1 to 2. • For *dscp1 ... dscp8*, specify the DSCP values that are mapped to a threshold ID. Enter up to eight DSCP values per command. Separate each value with a space. The DSCP range is 0 to 63.
Step 10	**switchport mode trunk**	Configure this port as a trunk port.
Step 11	**exit**	Return to global configuration mode.
Step 12	**mls qos map dscp-cos** *dscp-list* **to** *cos*	Modify the DSCP-to-CoS map. You can enter up to eight DSCP values separated by spaces in the DSCP-to-CoS map. For example, to map DSCP values 57 to 63 to CoS 5, enter: **mls qos map dscp-cos 57 58 59 60 61 62 63 to 5**
Step 13	**interface gigabitethernet0/5**	Enter interface configuration mode, and specify the egress interface to configure.
Step 14	**priority-queue out**	Enable the expedite queue.
Step 15	**wrr-queue bandwidth** *weight1 weight2 weight3 weight4*	Configure WRR weights to the egress queues to determine the ratio of the frequency at which packets are dropped. Separate each value with a space. The weight range is 0 to 65536. In this example, to configure the weights so that queue 4 is serviced more frequently than the other queues, enter: **wrr-queue bandwidth 1 2 3 4** Because the expedite queue is enabled, only the first three weights are used in the ratio calculation.
Step 16	**wrr-queue random-detect max-threshold** *queue-id threshold-percentage1 threshold-percentage2*	Enable WRED and assign two WRED threshold values to an egress queue of a Gigabit-capable Ethernet port. • For *queue-id*, the range is 1 to 4. • For *threshold-percentage1 threshold-percentage2*, the range is 1 to 100 percent. In this example, to configure the thresholds, enter: **wrr-queue random-detect max-threshold 1 20 100** **wrr-queue random-detect max-threshold 2 40 100** **wrr-queue random-detect max-threshold 3 60 100** **wrr-queue random-detect max-threshold 4 80 100**

	Command	Purpose
Step 17	**end**	Return to privileged EXEC mode.
Step 18	**show mls qos interface** and **show interfaces**	Verify your entries.
Step 19	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Configuring EtherChannel

This chapter describes how to configure EtherChannel on Layer 2 and Layer 3 interfaces. To configure Layer 3 interfaces, you must have the enhanced multilayer software image (EMI) installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This chapter consists of these sections:

- Understanding Port-Channel Interfaces, page 21-2
- Configuring EtherChannel, page 21-7
- Displaying EtherChannel and PAgP Status, page 21-16

Understanding EtherChannel

EtherChannel consists of individual Fast Ethernet or Gigabit Ethernet links bundled into a single logical link as shown in Figure 21-1. The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be the same speed, and all must be configured as either Layer 2 or Layer 3 interfaces.

Note The network device to which your switch is connected can impose its own limits on the number of interfaces in the EtherChannel. For Catalyst 3550 switches, the number of EtherChannels is limited to the number of ports of the same type.

If a link within an EtherChannel fails, traffic previously carried over that failed link changes to the remaining links within the EtherChannel. A trap is sent for a failure, identifying the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one link in an EtherChannel are blocked from returning on any other link of the EtherChannel.

Figure 21-1 Typical EtherChannel Configuration

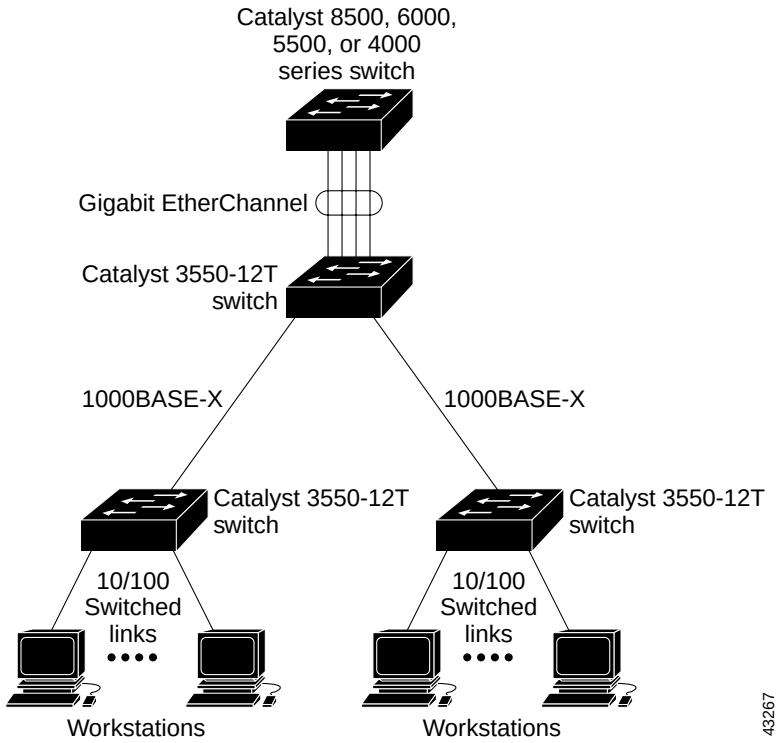

Understanding Port-Channel Interfaces

You create an EtherChannel for Layer 2 interfaces differently from Layer 3 interfaces. Both configurations involve logical interfaces.

- With Layer 3 interfaces, you manually create the logical interface by using the **interface port-channel** global configuration command.
- With Layer 2 interfaces, the logical interface is dynamically created.
- With both Layer 3 and 2 interfaces, you manually assign an interface to the EtherChannel by using the **channel-group** interface configuration command. This command binds the physical and logical ports together as shown in Figure 21-2.

Each EtherChannel has a logical port-channel interface numbered from 1 to 64. The channel groups are also numbered from 1 to 64.

Figure 21-2 Relationship of Physical Ports, Logical Port Channels, and Channel Groups

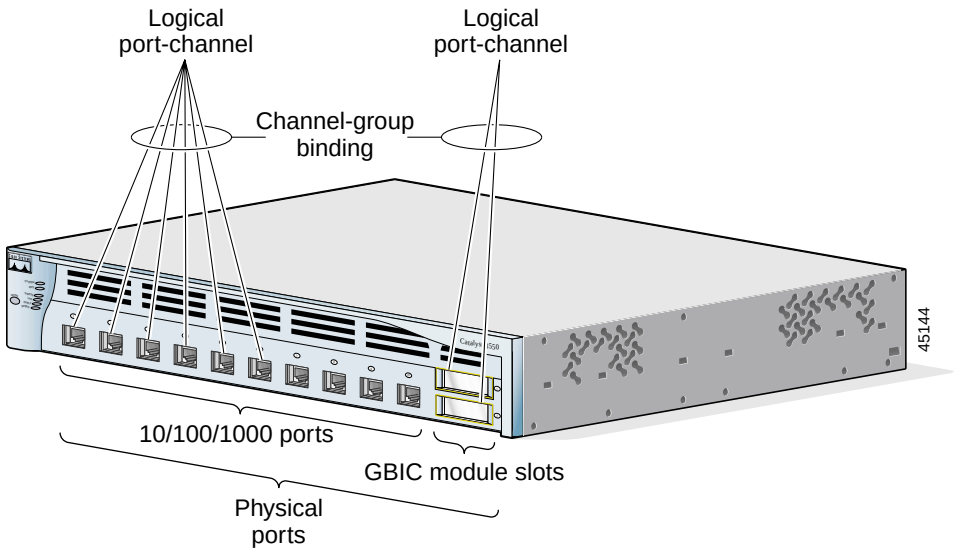

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface. Configuration changes applied to the physical interface affect only the interface where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface, for example, Spanning Tree Protocol (STP) commands or commands to configure a Layer 2 EtherChannel as a trunk.

Understanding the Port Aggregation Protocol

The Port Aggregation Protocol (PAgP) facilitates the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces. By using PAgP, the switch learns the identity of partners capable of supporting PAgP and learns the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port); these interfaces are grouped based on hardware, administrative, and port parameter constraints. For example, PAgP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP Modes

Table 21-1 shows the user-configurable EtherChannel modes for the **channel-group** interface configuration command: **on**, **auto**, and **desirable**. Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes; interfaces configured in the **on** mode do not exchange PAgP packets.

Table 21-1 EtherChannel Modes

Mode	Description
auto	Places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets and is the default.
desirable	Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets.
on	Forces the interface to channel without PAgP. With the **on** mode, a usable EtherChannel exists only when an interface group in the **on** mode is connected to another interface group in the **on** mode.

Both the **auto** and **desirable** modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- An interface in **desirable** mode can form an EtherChannel with another interface that is in **desirable** or **auto** mode.
- An interface in **auto** mode can form an EtherChannel with another interface in **desirable** mode.
- An interface in **auto** mode cannot form an EtherChannel with another interface that is also in **auto** mode because neither interface starts PAgP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.

Caution You should exercise care when setting the mode to **on** (manual configuration). All ports configured in the **on** mode are bundled in the same group and are forced to have similar characteristics. If the group is misconfigured, packet loss or STP loops might occur.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation by using the **non-silent** keyword. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

The silent mode is used when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational; however, the silent setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

Physical Learners and Aggregate-Port Learners

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate (logical) ports.

When a device and its partner are both aggregate-port learners, they learn the address on the logical port-channel. The device sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

PAgP cannot automatically detect when the partner device is a physical learner and the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device or source-based distribution by using the **pagp learn-method** interface configuration command. With source-based distribution, any given source MAC address is always sent on the same physical port. You can also configure a single interface within the group for all transmissions and use other interfaces for hot standby. The unused interfaces in the group can be swapped into operation in just a few seconds if the selected single interface loses hardware-signal detection. You can configure which interface is always selected for packet transmission by changing its priority by using the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.

PAgP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and Cisco Discovery Protocol (CDP) send and receive packets over the physical interfaces in the EtherChannel. Trunk ports send and receive PAgP protocol data units (PDUs) on the lowest numbered VLAN.

STP sends packets over the first interface in the EtherChannel.

The MAC address of a Layer 3 EtherChannel is the MAC address of the first interface in the port-channel.

PAgP sends and receives PAgP PDUs only from interfaces that are up and have PAgP enabled for the auto or desirable mode.

Understanding Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use either source-MAC or destination-MAC address forwarding.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel (and the MAC address learned by the switch does not change).

When source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address. Packets between two IP hosts always use the same port in the channel, and traffic between any other pair of hosts can use a different port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

In Figure 21-3, an EtherChannel of four workstations communicates with a router. Because the router is a single-MAC-address device, source-based forwarding on the switch EtherChannel ensures that the switch uses all available bandwidth to the router. The router is configured for destination-based forwarding because the large number of workstations ensures that the traffic is evenly distributed from the router EtherChannel.

Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel is going only to a single MAC address, using the destination-MAC address always chooses the same link in the channel; using source addresses or IP addresses might result in better load balancing.

Figure 21-3 Load Distribution and Forwarding Methods

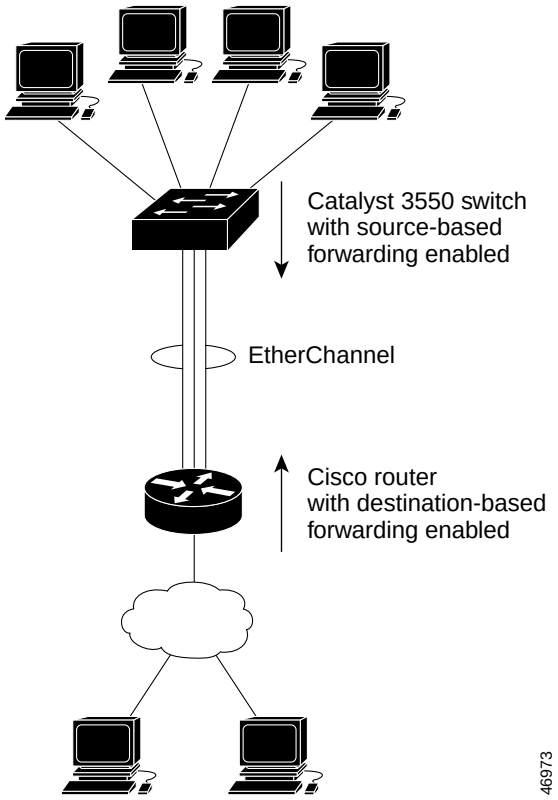

Configuring EtherChannel

This section describes these configurations for EtherChannel on Layer 2 and Layer 3 interfaces:

- Default EtherChannel Configuration, page 21-7
- EtherChannel Configuration Guidelines, page 21-8
- Configuring Layer 2 EtherChannels, page 21-9
- Configuring Layer 3 EtherChannels, page 21-11
- Configuring EtherChannel Load Balancing, page 21-13
- Configuring the PAgP Learn Method and Priority, page 21-14

Note Make sure that the interfaces are correctly configured (see the "EtherChannel Configuration Guidelines" section on page 21-8).

Note After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface, and configuration changes applied to the physical interface affect only the interface where you apply the configuration.

Default EtherChannel Configuration

Table 21-2 shows the default EtherChannel configuration.

Table 21-2 Default EtherChannel Configuration

Feature	Default Setting
Channel groups	None assigned.
Layer 3 port-channel logical interface	None defined.
PAgP mode	Auto and silent (The interface is in a passive negotiating state; it responds to PAgP packets it receives but does not start PAgP packet negotiation. PAgP is enabled only if a PAgP device is detected.)
PAgP learn method	Aggregate-port learning on all interfaces.
PAgP priority	128 on all interfaces.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet. Load distribution based on the source and destination IP address is also enabled for routed IP traffic.

EtherChannel Configuration Guidelines

If improperly configured, some EtherChannel interfaces are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- Each EtherChannel can have up to eight compatibly configured Ethernet interfaces.

  Note Do not configure a GigaStack GBIC port as part of an EtherChannel.

- Configure all interfaces in an EtherChannel to operate at the same speeds and duplex modes.
- Enable all interfaces in an EtherChannel. An interface in an EtherChannel disabled by using the **shutdown** interface configuration command is treated as a link failure, and its traffic is transferred to one of the remaining interfaces in the EtherChannel.
- When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, you must also make the changes to all ports in the group:
  - Allowed-VLAN list
  - STP path cost for each VLAN
  - STP port priority for each VLAN
  - STP Port Fast setting
- An EtherChannel does not form if one of the interfaces is a Switch Port Analyzer (SPAN) destination port.
- A port that belongs to an EtherChannel port group cannot be configured as a secure port.
- Before enabling 802.1X on the port, you must first remove it from the EtherChannel. If you try to enable 802.1X on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1X is not enabled. If you enable 802.1X on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
- For Layer 2 EtherChannels:
  - Assign all interfaces in the EtherChannel to the same VLAN, or configure them as trunks. Interfaces with different native VLANs cannot form an EtherChannel.
  - If you configure an EtherChannel from trunk interfaces, verify that the trunking mode (ISL or 802.1Q) is the same on all the trunks. Inconsistent trunk modes on EtherChannel interfaces can have unexpected results.
  - An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
  - Interfaces with different STP path costs can form an EtherChannel as long they are otherwise compatibly configured. Setting different STP path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.
- For Layer 3 EtherChannels, assign the Layer 3 address to the port-channel logical interface, not to the physical interfaces in the channel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by configuring the Ethernet interfaces with the **channel-group** interface configuration command, which creates the port-channel logical interface.

Note Layer 2 interfaces must be connected and functioning for IOS to create port-channel interfaces for Layer 2 EtherChannels.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. Up to eight interfaces of the same type and speed can be configured for the same group.
Step 3	**switchport mode** {**access** | **trunk**} **switchport access vlan** *vlan-id*	Assign all interfaces as static-access ports in the same VLAN, or configure them as trunks. If you configure the interface as a static-access port, assign it to only one VLAN. The range is 1 to 1005.

	Command	Purpose
Step 4	**channel-group** *channel-group-number* **mode** {**auto** [**non-silent**] | **desirable** [**non-silent**] | **on**}	Assign the interface to a channel group, and specify the PAgP mode. The default mode is **auto silent**. For *channel-group-number*, the range is 1 to 64. Each EtherChannel can have of up to eight compatibly configured Ethernet interfaces. For **mode**, select one of these keywords: • **auto**—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • **desirable**—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • **on**—Forces the interface to channel without PAgP. With the **on** mode, a usable EtherChannel exists only when an interface group in the **on** mode is connected to another interface group in the **on** mode. • **non-silent**—If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation. You can configure an interface with the **non-silent** keyword for use with the **auto** or **desirable** mode. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers; this setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. For information on compatible PAgP modes for the switch and its partner, see the "PAgP Modes" section on page 21-4.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to assign Gigabit Ethernet interfaces 0/4 and 0/5 as static-access ports in VLAN 10 to channel 5 with PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/4 -5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable
Switch(config-if-range)# end
```

Note For information about the **range** keyword, see the "Configuring a Range of Interfaces" section on page 8-9.

Configuring Layer 3 EtherChannels

To configure Layer 3 EtherChannels, you create the port-channel logical interface and then put the Ethernet interfaces into the port-channel as described in the next two sections.

Creating Port-Channel Logical Interfaces

Note To move an IP address from a physical interface to an EtherChannel, you must delete the IP address from the physical interface before configuring it on the port-channel interface.

Beginning in privileged EXEC mode, follow these steps to create a port-channel interface for a Layer 3 EtherChannel:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface port-channel** *port-channel-number*	Enter interface configuration mode, and create the port-channel logical interface. For *port-channel-number*, the range is 1 to 64.
Step 3	**no switchport**	Put the interface into Layer 3 mode. You must have the EMI installed to use this command.
Step 4	**ip address** *ip-address mask*	Assign an IP address and subnet mask to the EtherChannel.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show etherchannel** *channel-group-number* **detail**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.
Step 8		Assign an Ethernet interface to the Layer 3 EtherChannel. For more information, see the "Configuring the Physical Interfaces" section on page 21-12.

To remove the port-channel, use the **no interface port-channel** *port-channel-number* global configuration command.

This example shows how to create the logical port channel (5) and assign 172.10.20.10 as its IP address:

```
Switch# configure terminal
Switch(config)# interface port-channel 5
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.10.20.10 255.255.255.0
Switch(config-if)# end
```

Configuring the Physical Interfaces

Beginning in privileged EXEC mode, follow these steps to assign an Ethernet interface to a Layer 3 EtherChannel:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. Up to eight interfaces of the same type and speed can be configured for the same group.
Step 3	**no ip address**	Ensure that there is no IP address assigned to the physical interface.
Step 4	**channel-group** *channel-group-number* **mode** {**auto** [**non-silent**] | **desirable** [**non-silent**] | **on**}	Assign the interface to a channel group, and specify the PAgP mode (the default mode is **auto silent**). For *channel-group-number*, the range is 1 to 64. This number must be the same as the *port-channel-number* (logical port) configured in the "Creating Port-Channel Logical Interfaces" section on page 21-11. Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. For **mode**, select one of these keywords: • **auto**—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • **desirable**—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • **on**—Forces the interface to channel without PAgP. With the **on** mode, a usable EtherChannel exists only when an interface group in the **on** mode is connected to another interface group in the **on** mode. • **non-silent**—If your switch is connected to a partner that is PAgP capable, you can configure the switch interface for nonsilent operation. You can configure an interface with the **non-silent** keyword for use with the **auto** or **desirable** mode. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers; this setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. For information on compatible PAgP modes for the switch and its partner, see the "PAgP Modes" section on page 21-4.

	Command	Purpose
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to assign Gigabit Ethernet interfaces 0/4 and 0/5 to channel 5 with PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/4 -5
Switch(config-if-range)# no ip address
Switch(config-if-range)# channel-group 5 mode desirable
Switch(config-if-range)# end
```

Note For information about the **range** keyword, see the "Configuring a Range of Interfaces" section on page 8-9.

Configuring EtherChannel Load Balancing

This section describes how to configure EtherChannel load balancing by using source-based or destination-based forwarding methods. For more information, see the "Understanding Load Balancing and Forwarding Methods" section on page 21-5.

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**port-channel load-balance** {**dst-mac** | **src-mac**}	Configure an EtherChannel load-balancing method. The default is **src-mac**. Select one of these keywords to determine the load-distribution method: • **dst-mac**—Load distribution is based on the destination-host MAC address of the incoming packet. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel. • **src-mac**—Load distribution is based on the source-MAC address of the incoming packet. Packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel. When **src-mac** is used, load distribution based on the source and destination IP address is also enabled. For all IP traffic being routed, the switch chooses a port for transmission based on the source and destination IP address. Packets between two IP hosts always use the same port for packet transmission, but packets between any other pair of hosts might use a different transmission port.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show etherchannel load-balance**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return EtherChannel load balancing to the default configuration, use the **no port-channel load-balance** global configuration command.

Configuring the PAgP Learn Method and Priority

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate ports.

For compatibility with Catalyst 1900 series switches, configure the PAgP learning method on the Catalyst 3550 switches to learn source-MAC addresses on the physical port. The switch then sends packets to the Catalyst 1900 switch using the same interface in the EtherChannel from which it learned the source address.

Note The Catalyst 3550 supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the CLI. The **pagp learn-method** command and the **pagp port-priority** command have no effect on the switch hardware, but they are required for PAgP interoperability with devices that only support address learning by physical ports, such as the Catalyst 1900 switch.

When the link partner to the Catalyst 3550 switch is a physical learner, we recommend that you configure the switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command and to set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. Use the **pagp learn-method** command only in this situation.

Beginning in privileged EXEC mode, follow these steps to configure your switch as a PAgP physical-port learner and to adjust the priority so that the same port in the bundle is selected for sending packets:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface for transmission.
Step 3	**pagp learn-method physical-port**	Select the PAgP learning method. By default, **aggregation-port** learning is selected, which means the switch sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives. Select **physical-port** to connect with another switch that is a physical learner. Make sure to configure the **port-channel load-balance** global configuration command to **src-mac** as described in the "Configuring EtherChannel Load Balancing" section on page 21-13. The learning method must be configured the same at both ends of the link.
Step 4	**pagp port-priority** *priority*	Assign a priority so that the selected interface is chosen for packet transmission. For *priority*, the range is 0 to 255. The default is 128. The higher the priority, the more likely that the interface will be used for PAgP transmission.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config** or **show pagp** *channel-group-number* **internal**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return the priority to its default setting, use the **no pagp port-priority** interface configuration command. To return the learning method to its default setting, use the **no pagp learn-method** interface configuration command.

Displaying EtherChannel and PAgP Status

You can use the privileged EXEC commands described in Table 21-3 to display EtherChannel and PAgP status information:

Table 21-3 Commands for Displaying EtherChannel and PAgP Status

Command	Description
show etherchannel [*channel-group-number*] {**brief** | **detail** | **load-balance** | **port** | **port-channel** | **summary**}	Displays EtherChannel information in a brief, detailed, and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, and port-channel information.
show pagp [*channel-group-number*] {**counters** | **internal** | **neighbor**}[1]	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.

1. You can clear PAgP channel-group information and traffic filters by using the **clear pagp** {*channel-group-number* | **counters**} privileged EXEC command.

For detailed information about the fields in the displays, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

CHAPTER 22

Configuring IP Unicast Routing

This chapter describes how to configure IP unicast routing on your multilayer switch. To use this feature, you must have the enhanced multilayer software image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Note For more detailed IP unicast configuration information, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*. For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding Routing, page 22-2
- Steps for Configuring Routing, page 22-3
- Configuring IP Addressing, page 22-4
- Enabling IP Routing, page 22-24
- Configuring RIP, page 22-25
- Configuring IGRP, page 22-30
- Configuring OSPF, page 22-35
- Configuring EIGRP, page 22-46
- Configuring Protocol-Independent Features, page 22-53
- Monitoring and Maintaining the IP Network, page 22-64

Note When configuring routing parameters on the switch, to allocate system resources to maximize the number of unicast routes allowed, you can use the **sdm prefer routing** global configuration command to set the Switch Database Management feature to the routing template. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

Understanding Routing

Network devices in different VLANs cannot communicate with one another without a Layer 3 device (router) to route traffic between the VLANs. Routers can perform routing in three different ways:

- By using default routing
- By using preprogrammed static routes for the traffic
- By dynamically calculating routes by using a routing protocol

Default routing refers to sending traffic with a destination unknown to the router to a default outlet or destination.

Static routing forwards packets from predetermined ports through a single path into and out of a network. Static routing is secure and uses little bandwidth, but does not automatically respond to changes in the network, such as link failures, and therefore, might result in unreachable destinations. As networks grow, static routing becomes a labor-intensive liability.

Dynamic routing protocols are used by routers to dynamically calculate the best route for forwarding traffic. There are two types of dynamic routing protocols:

- Routers using distance-vector protocols maintain routing tables with distance values of networked resources, and periodically pass these tables to their neighbors. Distance-vector protocols use one or a series of metrics for calculating the best routes. These protocols are easy to configure and use.
- Routers using link-state protocols maintain a complex database of network topology, based on the exchange of link-state advertisements (LSAs) between routers. LSAs are triggered by an event in the network, which speeds up the convergence time or time required to respond to these changes. Link-state protocols respond quickly to topology changes, but require greater bandwidth and more resources than distance-vector protocols.

Distance-vector protocols supported by the Catalyst 3550 switch are Routing Information Protocol (RIP), which uses a single distance metric (cost) to determine the best path, and Interior Gateway Routing Protocol (IGRP), which uses a series of metrics. The switch also supports the Open Shortest Path First (OSPF) link-state protocol and Enhanced IGRP (EIGRP), which adds some link-state routing features to traditional IGRP to improve efficiency.

In some network environments, VLANs are associated with individual networks or subnetworks. In an IP network, each subnetwork is mapped to an individual VLAN. Configuring VLANs helps control the size of the broadcast domain and keeps local traffic local. However, when an end station in one VLAN needs to communicate with an end station in another VLAN, inter-VLAN communication is required. This communication is supported by inter-VLAN routing. You configure one or more routers to route traffic to the appropriate destination VLAN.

Figure 22-1 shows a basic routing topology. Switch A is in VLAN 10, and Switch B is in VLAN 20. The router has an interface in each VLAN.

Figure 22-1 Routing Topology Example

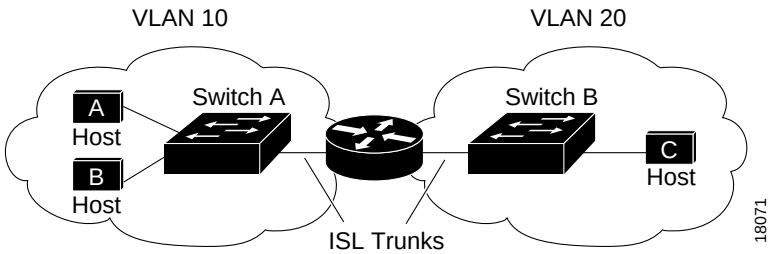

When Host A in VLAN 10 needs to communicate with Host B in VLAN 10, it sends a packet addressed to that host. Switch A forwards the packet directly to Host B, without sending it to the router.

When Host A sends a packet to Host C in VLAN 20, Switch A forwards the packet to the router, which receives the traffic on the VLAN 10 interface. The router checks the routing table, determines the correct outgoing interface, and forwards the packet on the VLAN 20 interface to Switch B. Switch B receives the packet and forwards it to Host C.

Steps for Configuring Routing

By default, IP routing is disabled on the Catalyst 3550 switch, and you must enable it before routing can take place. For detailed IP routing configuration information, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*.

In the following procedures, the specified interface must be one of these Layer 3 interfaces:

- A routed port: a physical port configured as a Layer 3 port by using the **no switchport** interface configuration command.
- A switch virtual interface (SVI): a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface.
- An EtherChannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel** *port-channel-number* global configuration command and binding the Ethernet interface into the channel group. For more information, see the "Configuring Layer 3 EtherChannels" section on page 21-11.

Note A Layer 3 switch can have an IP address assigned to each routed port and SVI. The number of routed ports and SVIs that you can configure is not limited by software. However, the interrelationship between this number and the number and volume of features being implemented might have an impact on CPU utilization because of hardware limitations. For more information about feature combinations, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

All Layer 3 interfaces must have IP addresses assigned to them. See the "Assigning IP Addresses to Network Interfaces" section on page 22-5.

Configuring routing consists of several main procedures:

- To support VLAN interfaces, create and configure VLANs on the switch, and assign VLAN membership to Layer 2 interfaces. For more information, see Chapter 9, "Creating and Maintaining VLANs."
- Configure Layer 3 interfaces.
- Enable IP routing on the switch.
- Assign IP addresses to the Layer 3 interfaces.
- Enable selected routing protocols on the switch.
- Configure routing protocol parameters (optional).

Configuring IP Addressing

A required task for configuring IP routing is to assign IP addresses to Layer 3 network interfaces to enable the interfaces and allow communication with the hosts on those interfaces that use IP. These sections describe how to configure various IP addressing features. Assigning IP addresses to the interface is required; the other procedures are optional.

- Default Addressing Configuration, page 22-4
- Assigning IP Addresses to Network Interfaces, page 22-5
- Configuring Address Resolution Methods, page 22-10
- Routing Assistance When IP Routing is Disabled, page 22-14
- Configuring Broadcast Packet Handling, page 22-17
- Monitoring and Maintaining IP Addressing, page 22-21

Default Addressing Configuration

Table 22-1 shows the default addressing configuration.

Table 22-1 Default Addressing Configuration

Feature	Default Setting
IP address	None defined.
ARP	No permanent entries in the Address Resolution Protocol (ARP) cache. Encapsulation: Standard Ethernet-style ARP. Timeout: 14400 seconds (4 hours).
IP broadcast address	255.255.255.255 (all ones).
IP classless routing	Enabled.
IP default gateway	Disabled.
IP directed broadcast	Disabled (all IP directed broadcasts are dropped).
IP domain	Domain list: No domain names defined. Domain lookup: Enabled. Domain name: Enabled.
IP forward-protocol	If a helper address is defined or User Datagram Protocol (UDP) flooding is configured, UDP forwarding is enabled on default ports. Any-local-broadcast: Disabled. Spanning Tree Protocol (STP): Disabled. Turbo-flood: Disabled.
IP helper address	Disabled.
IP host	Disabled.

Table 22-1 Default Addressing Configuration (continued)

Feature	Default Setting
IRDP	Disabled. Defaults when enabled: • Broadcast IRDP advertisements. • Maximum interval between advertisements: 600 seconds. • Minimum interval between advertisements: 0.75 times max interval • Preference: 0.
IP proxy ARP	Enabled.
IP routing	Disabled.
IP subnet-zero	Disabled.

Assigning IP Addresses to Network Interfaces

An IP address identifies a location to which IP packets can be sent. Some IP addresses are reserved for special uses and cannot be used for host, subnet, or network addresses. Table 22-2 lists ranges of IP addresses and shows which are reserved and which are available for use. RFC 1166, “Internet Numbers,” contains the official description of IP addresses.

Table 22-2 Reserved and Available IP Addresses

Class	Address or Range	Status
A	0.0.0.0 1.0.0.0 to 126.0.0.0 127.0.0.0	Reserved Available Reserved
B	128.0.0.0 to 191.254.0.0 191.255.0.0	Available Reserved
C	192.0.0.0 192.0.1.0 to 223.255.254 223.255.255.0	Reserved Available Reserved
D	224.0.0.0 to 239.255.255.255	Multicast group addresses
E	240.0.0.0 to 255.255.255.254 255.255.255.255	Reserved Broadcast

An interface can have one primary IP address. A mask identifies the bits that denote the network number in an IP address. When you use the mask to subnet a network, the mask is referred to as a subnet mask. To receive an assigned network number, contact your Internet service provider.

Beginning in privileged EXEC mode, follow these steps to assign an IP address and a network mask to a Layer 3 interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**no switchport**	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 4	**ip address** *ip-address subnet-mask*	Configure the IP address and IP subnet mask.
Step 5	**no shutdown**	Enable the interface.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show interfaces** [*interface-id*] **show ip interface** [*interface-id*] **show running-config interface** [*interface-id*]	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip address** interface configuration command to remove an IP address or to disable IP processing.

This example shows how to configure an IP address on Gigabit Ethernet interface 0/10:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface gigabitethernet0/10
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.1.2.3 255.255.0.0
Switch(config-if)# no shutdown
Switch(config-if)# end
Switch#
```

This is an example of output from the **show interfaces** privileged EXEC command for Gigabit Ethernet interface 0/10, displaying the interface IP address configuration and status:

```
Switch# show interfaces gigabitethernet0/10
GigabitEthernet0/10 is up, line protocol is down
  Hardware is Gigabit Ethernet, address is 0002.4b29.7100 (bia 0002.4b29.7100)
  Internet address is 10.1.2.3/16
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:42, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
```

```
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 2 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

This is an example of output from the **show ip interface** privileged EXEC command for Gigabit Ethernet interface 0/10, displaying the detailed IP configuration and status:

```
Switch# show ip interface gigabitethernet0/10
GigabitEthernet0/10 is up, line protocol is down
  Internet address is 10.1.2.3/16
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect exclude is disabled
  BGP Policy Mapping is disabled
```

This is an example of output from the **show running-config** privileged EXEC command for of Gigabit Ethernet interface 0/10 to display the interface IP address configuration:

```
Switch# show running-config interface gigabitethernet0/10
Building configuration...

Current configuration : 189 bytes
!
interface GigabitEthernet0/10
 description CubeB
 no switchport
 ip address 10.1.2.3 255.255.0.0
 ip access-group 23 in
 ip access-group stan1 out
 ip pim sparse-dense-mode
 ip cgmp
end
```

Use of Subnet Zero

Subnetting with a subnet address of zero is strongly discouraged because of the problems that can arise if a network and a subnet have the same addresses. For example, if network 131.108.0.0 is subnetted as 255.255.255.0, subnet zero would be written as 131.108.0.0, which is the same as the network address.

You can use the all ones subnet (131.108.255.0) and even though it is discouraged, you can enable the use of subnet zero if you need the entire subnet space for your IP address.

Beginning in privileged EXEC mode, follow these steps to enable subnet zero:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip subnet-zero**	Enable the use of subnet zero for interface addresses and routing updates.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

Use the **no ip subnet-zero** global configuration command to restore the default and disable the use of subnet zero.

This is an example of partial output from the **show running-config** privileged EXEC command used to verify IP subnet zero setting.

```
Switch# show running-config
Building configuration...

Current configuration : 7454 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Perdido1
!
!
<output truncated>
ip subnet-zero
ip routing
no ip domain-lookup
ip domain-name a,b,c
ip name-server 12.10.13.14

<output truncated>
```

Classless Routing

By default, classless routing behavior is enabled on the switch when it is configured to route. With classless routing, if a router receives packets for a subnet of a network with no default route, the router forwards the packet to the best supernet route. A *supernet* consists of contiguous blocks of Class C address spaces used to simulate a single, larger address space and is designed to relieve the pressure on the rapidly depleting Class B address space.

In Figure 22-2, classless routing is enabled. When the host sends a packet to 120.20.4.1, instead of discarding the packet, the router forwards it to the best supernet route. If you disable classless routing and a router receives packets destined for a subnet of a network with no network default route, the router discards the packet.

Figure 22-2 IP Classless Routing

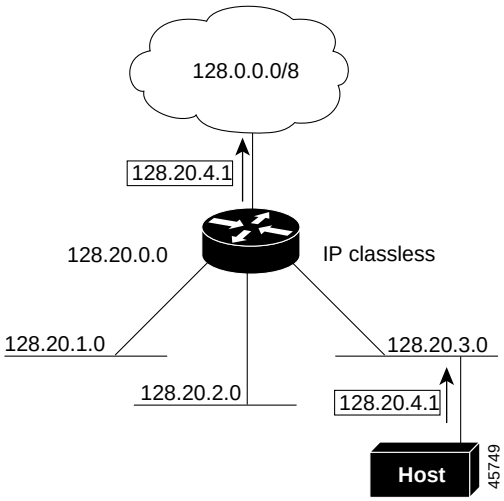

In Figure 22-3, the router in network 128.20.0.0 is connected to subnets 128.20.1.0, 128.20.2.0, and 128.20.3.0. If the host sends a packet to 120.20.4.1, because there is no network default route, the router discards the packet.

Figure 22-3 No IP Classless Routing

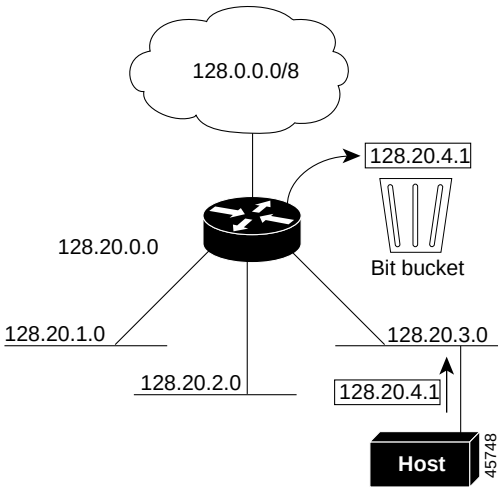

To prevent the switch from forwarding packets destined for unrecognized subnets to the best supernet route possible, you can disable classless routing behavior.

Beginning in privileged EXEC mode, follow these steps to disable classless routing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no ip classless**	Disable classless routing behavior.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To restore the default and have the switch forward packets destined for a subnet of a network with no network default route to the best supernet route possible, use the **ip classless** global configuration command.

Configuring Address Resolution Methods

You can control interface-specific handling of IP by using address resolution. A device using IP can have both a local address or MAC address, which uniquely defines the device on its local segment or LAN, and a network address, which identifies the network to which the device belongs. The local address or MAC address is known as a data link address because it is contained in the data link layer (Layer 2) section of the packet header and is read by data link (Layer 2) devices. To communicate with a device on Ethernet, the software must determine the MAC address of the device. The process of determining the MAC address from an IP address is called *address resolution*. The process of determining the IP address from the MAC address is called *reverse address resolution*.

The switch can use these forms of address resolution:

- Address Resolution Protocol (ARP) is used to associate IP address with MAC addresses. Taking an IP address as input, ARP determines the associated MAC address and then stores the IP address/MAC address association in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests or replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP).
- Proxy ARP helps hosts with no routing tables determine the MAC addresses of hosts on other networks or subnets. If the switch (router) receives an ARP request for a host that is not on the same interface as the ARP request sender, and if the router has all of its routes to the host through other interfaces, it generates a proxy ARP packet giving its own local data link address. The host that sent the ARP request then sends its packets to the router, which forwards them to the intended host.

Catalyst 3550 switches also use the Reverse Address Resolution Protocol (RARP), which functions the same as ARP does, except that the RARP packets request an IP address instead of a local MAC address. Using RARP requires a RARP server on the same network segment as the router interface. Use the **ip rarp-server** *address* interface configuration command to identify the server.

For more information on RARP, refer to the *Cisco IOS Configuration Fundamentals Configuration Guide for Release 12.1*.

You can perform these tasks to configure address resolution:

- Define a Static ARP Cache, page 22-11
- Set ARP Encapsulation, page 22-12
- Enable Proxy ARP, page 22-13

Define a Static ARP Cache

ARP and other address resolution protocols provide dynamic mapping between IP addresses and MAC addresses. Because most hosts support dynamic address resolution, you usually do not need to specify static ARP cache entries. If you must define a static ARP cache entry, you can do so globally, which installs a permanent entry in the ARP cache that the switch uses to translate IP addresses into MAC addresses. Optionally, you can also specify that the switch respond to ARP requests as if it were the owner of the specified IP address. If you do not want the ARP entry to be permanent, you can specify a timeout period for the ARP entry.

Beginning in privileged EXEC mode, follow these steps to provide static mapping between IP addresses and MAC addresses:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**arp** *ip-address hardware-address type*	Globally associate an IP address with a MAC (hardware) address in the ARP cache, and specify encapsulation type as one of these: • **arpa**—ARP encapsulation for Ethernet interfaces • **snap**—Subnetwork Address Protocol encapsulation for Token Ring and FDDI interfaces • **sap**—HP's ARP type
Step 3	**arp** *ip-address hardware-address type* [**alias**]	(Optional) Specify that the switch respond to ARP requests as if it were the owner of the specified IP address.
Step 4	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 5	**arp timeout** *seconds*	(Optional) Set the length of time an ARP cache entry will stay in the cache. The default is 14400 seconds (4 hours). The range is 0 to 2147483 seconds.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show interfaces** [*interface-id*]	Verify the type of ARP and the timeout value used on all interfaces or a specific interface.
Step 8	**show arp** or **show ip arp**	View the contents of the ARP cache.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove an entry from the ARP cache, use the **no arp** *ip-address hardware-address type* global configuration command. To remove all nonstatic entries from the ARP cache, use the **clear arp-cache** privileged EXEC command.

This is an example of output from the **show arp** privileged EXEC command.

```
Switch# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.2.3                -   0002.4b29.2e00  ARPA   GigabitEthernet0/10
Internet  172.20.136.9          120   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.250.42         149   0030.19c6.54e1  ARPA   Vlan1
Internet  120.20.30.1             -   0002.4b29.2e00  ARPA   Vlan27
Internet  172.20.139.152        101   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.139.130        205   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.141.225        186   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.135.204        169   0002.4b29.4400  ARPA   Vlan1
Internet  172.20.135.202          -   0002.4b29.2e00  ARPA   Vlan1
Internet  172.20.135.197        172   0002.4b28.ce80  ARPA   Vlan1
Internet  172.20.135.196        156   0002.4b28.ce00  ARPA   Vlan1
```

Note For the Catalyst 3550 switch, the output from the **show arp** privileged EXEC command and the **show ip arp** privileged EXEC command are usually the same.

Set ARP Encapsulation

By default, Ethernet ARP encapsulation (represented by the **arpa** keyword) is enabled on an IP interface. You can change the encapsulation methods to SNAP if required by your network.

Beginning in privileged EXEC mode, follow these steps to specify the ARP encapsulation type:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**arp** {**arpa** | **snap**}	Specify the ARP encapsulation method: • **arpa**—Address Resolution Protocol • **snap**—Subnetwork Address Protocol
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show interfaces** [*interface-id*]	Verify ARP encapsulation configuration on all interfaces or the specified interface.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable an encapsulation type, use the **no arp arpa** or **no arp snap** interface configuration command.

This is an example of output from the **show interfaces** *interface-id* privileged EXEC command displaying ARP encapsulation.

```
Switch# show interfaces gigabitethernet0/10
GigabitEthernet0/10 is up, line protocol is up
  Hardware is Gigabit Ethernet, address is 0002.4b29.2e00 (bia 0002
  Internet address is 40.5.121.10/24
  MTU 1500 bytes, BW 100000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:04, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     30745 packets output, 3432096 bytes, 0 underruns
     0 output errors, 0 collisions, 6 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

Enable Proxy ARP

By default, the switch uses proxy ARP to help hosts determine MAC addresses of hosts on other networks or subnets.

Beginning in privileged EXEC mode, follow these steps to enable proxy ARP if it has been disabled:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip proxy-arp**	Enable proxy ARP on the interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip interface** [*interface-id*]	Verify the configuration on the interface or all interfaces.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable proxy ARP on the interface, use the **no ip proxy-arp** interface configuration command.

This is an example of output form the **show ip interface** privileged EXEC command for Gigabit Ethernet interface 0/3, where proxy ARP is enabled.

```
Switch# show ip interface gigabitethernet0/3
GigabitEthernet0/3 is up, line protocol is down
  Internet address is 10.1.3.59/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.1 224.0.0.2
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast, CEF
  Router Discovery is enabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect exclude is disabled
  BGP Policy Mapping is disabled
```

Routing Assistance When IP Routing is Disabled

These mechanisms allow the switch to learn about routes to other networks when it does not have IP routing enabled:

- Proxy ARP, page 22-14
- Default Gateway, page 22-15
- ICMP Router Discovery Protocol (IRDP), page 22-15

Proxy ARP

Proxy ARP, the most common method for learning about other routes, enables an Ethernet host with no routing information to communicate with hosts on other networks or subnets. The host assumes that all hosts are on the same local Ethernet and that they can use ARP to determine their MAC addresses. If a switch receives an ARP request for a host that is not on the same network as the sender, the switch evaluates whether it has the best route to that host. If it does, it sends an ARP reply packet with its own

Ethernet MAC address, and the host that sent the request sends the packet to the switch, which forwards it to the intended host. Proxy ARP treats all networks as if they are local and performs ARP requests for every IP address.

Proxy ARP is enabled by default. To enable it after it has been disabled, see the "Enable Proxy ARP" section on page 22-13. Proxy ARP works as long as other routers support it.

Default Gateway

Another method for locating routes is to define a default router or default gateway. All nonlocal packets are sent to this router, which either routes them appropriately or sends an IP Control Message Protocol (ICMP) redirect message back, defining which local router the host should use. The switch caches the redirect messages and forwards each packet as efficiently as possible. A limitation of this method is that there is no means of detecting when the default router has gone down or is unavailable.

Beginning in privileged EXEC mode, follow these steps to define a default gateway (router) when IP routing is disabled:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip default-gateway** *ip-address*	Set up a default gateway (router).
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip redirects**	Display the address of the default gateway router to verify the setting.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip default-gateway** global configuration command to disable this function.

This example shows how to set and verify a default gateway:

```
Switch(config)# ip default-gateway 10.1.5.59
Switch(config)# end
Switch# show ip redirect
Default gateway is 10.1.5.59

Host              Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

ICMP Router Discovery Protocol (IRDP)

Router discovery allows the switch to dynamically learn about routes to other networks using IRDP. IRDP allows hosts to locate routers. When operating as a client, the switch generates router discovery packets. When operating as a host, the switch receives router discovery packets. The switch can also listen to Routing Information Protocol (RIP) and Interior Gateway Routing Protocol (IGRP) routing updates and use this information to infer locations of routers. The switch does not actually store the routing tables sent by routing devices; it merely keeps track of which systems are sending the data. The advantage of using IRDP is that it allows each router to specify both a priority and the time after which a device is assumed to be down if no further packets are received.

Each device discovered becomes a candidate for the default router, and a new highest-priority router is selected when a higher priority router is discovered, when the current default router is declared down, or when a TCP connection is about to time out because of excessive retransmissions.

The only required task for IRDP routing on an interface is to enable IRDP processing on that interface. When enabled, the default parameters apply. You can optionally change any of these parameters.

Beginning in privileged EXEC mode, follow these steps to enable and configure IRDP on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip irdp**	Enable IRDP processing on the interface.
Step 4	**ip irdp multicast**	(Optional) Send IRDP advertisements to the multicast address (224.0.0.1) instead of IP broadcasts. **Note** This command allows for compatibility with Sun Microsystems Solaris, which requires IRDP packets to be sent out as multicasts. Many implementations cannot receive these multicasts; ensure end-host ability before using this command.
Step 5	**ip irdp holdtime** *seconds*	(Optional) Set the IRDP period for which advertisements are valid. The default is three times the **maxadvertinterval** value. It must be greater than **maxadvertinterval** and cannot be greater than 9000 seconds. If you change the **maxadvertinterval** value, this value also changes.
Step 6	**ip irdp maxadvertinterval** *seconds*	(Optional) Set the IRDP maximum interval between advertisements. The default is 600 seconds.
Step 7	**ip irdp minadvertinterval** *seconds*	(Optional) Set the IRDP minimum interval between advertisements. The default is 0.75 times the **maxadvertinterval**. If you change the **maxadvertinterval**, this value changes to the new default (0.75 of **maxadvertinterval**).
Step 8	**ip irdp preference** *number*	(Optional) Set a device IRDP preference level. The allowed range is -2^{31} to 2^{31}. The default is 0. A higher value increases the router preference level.
Step 9	**ip irdp address** *address* [*number*]	(Optional) Specify an IRDP address and preference to proxy-advertise.
Step 10	**end**	Return to privileged EXEC mode.
Step 11	**show ip irdp**	Verify settings by displaying IRDP values.
Step 12	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

If you change the **maxadvertinterval** value, the **holdtime** and **minadvertinterval** values also change, so it is important to first change the **maxadvertinterval** value, before manually changing either the **holdtime** or **minadvertinterval** values.

Use the **no ip irdp** interface configuration command to disable IRDP routing.

This is an example of output from the **show ip irdp** privileged EXEC command that shows that IRDP routing is enabled on Gigabit Ethernet interface 0/3:

```
Switch# show ip irdp
Vlan1 has router discovery disabled

Vlan2 has router discovery disabled

GigabitEthernet0/1 has router discovery disabled

GigabitEthernet0/2 has router discovery disabled
```

```
GigabitEthernet0/3 has router discovery enabled

Advertisements will occur between every 450 and 600 seconds.
Advertisements are sent with broadcasts.
Advertisements are valid for 1800 seconds.
Default preference will be 0.
GigabitEthernet0/4 has router discovery disabled

Port-channel1 has router discovery disabled
```

Configuring Broadcast Packet Handling

After configuring an IP interface address, you can choose to enable routing and configure one or more routing protocols, or you can configure the way the switch responds to network broadcasts. A broadcast is a data packet destined for all hosts on a physical network. The switch supports two kinds of broadcasting:

- A directed broadcast packet is sent to a specific network or series of networks. A directed broadcast address includes the network or subnet fields.
- A flooded broadcast packet is sent to every network.

Note You can also limit broadcast, unicast, and multicast traffic on Layer 2 interfaces by using the **switchport broadcast**, **switchport unicast**, and **switchport multicast** interface configuration commands. For more information, see Chapter 12, "Configuring Port-Based Traffic Control."

Routers provide some protection from broadcast storms by limiting their extent to the local cable. Bridges (including intelligent bridges), because they are Layer 2 devices, forward broadcasts to all network segments, thus propagating broadcast storms. The best solution to the broadcast storm problem is to use a single broadcast address scheme on a network. In most modern IP implementations, you can set the address to be used as the broadcast address. Many implementations, including the one in the Catalyst 3550 switch, support several addressing schemes for forwarding broadcast messages.

Perform the tasks in these sections to enable these schemes:

- Enabling Directed Broadcast-to-Physical Broadcast Translation, page 22-17
- Forwarding UDP Broadcast Packets and Protocols, page 22-18
- Establishing an IP Broadcast Address, page 22-20
- Flooding IP Broadcasts, page 22-20

Enabling Directed Broadcast-to-Physical Broadcast Translation

By default, IP directed broadcasts are dropped; they are not forwarded. Dropping IP-directed broadcasts makes routers less susceptible to denial-of-service attacks.

You can enable forwarding of IP-directed broadcasts on an interface where the broadcast becomes a physical (MAC-layer) broadcast. Only those protocols configured by using the **ip forward-protocol** global configuration command are forwarded.

You can specify an access list to control which broadcasts are forwarded. When an access list is specified, only those IP packets permitted by the access list are eligible to be translated from directed broadcasts to physical broadcasts. For more information on access lists, see Chapter 19, "Configuring Network Security with ACLs."

Beginning in privileged EXEC mode, follow these steps to enable forwarding of IP-directed broadcasts on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 3	**ip directed-broadcast** [*access-list-number*]	Enable directed broadcast-to-physical broadcast translation on the interface. You can include an access list to control which broadcasts are forwarded. When an access list is specified, only IP packets permitted by the access list are eligible to be translated.
Step 4	**exit**	Return to global configuration mode.
Step 5	**ip forward-protocol** {**udp** [*port*] | **nd** | **sdns**}	Specify which protocols and ports the router forwards when forwarding broadcast packets. • **udp**—Forward UPD datagrams. *port*: (Optional) Destination port that controls which UDP services are forwarded. • **nd**—Forward ND datagrams. • **sdns**—Forward SDNS datagrams
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show ip interface** [*interface-id*] or **show running-config**	Verify the configuration on the interface or all interfaces.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip directed-broadcast** interface configuration command to disable translation of directed broadcast to physical broadcasts. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

Forwarding UDP Broadcast Packets and Protocols

User Datagram Protocol (UDP) is an IP host-to-host layer protocol, as is TCP. UDP provides a low-overhead, connectionless session between two end systems and does not provide for acknowledgment of received datagrams. Network hosts occasionally use UDP broadcasts to determine address, configuration, and name information. If such a host is on a network segment that does not include a server, UDP broadcasts are normally not forwarded. You can remedy this situation by configuring an interface on a router to forward certain classes of broadcasts to a helper address. You can use more than one helper address per interface.

You can specify a UDP destination port to control which UDP services are forwarded. You can specify multiple UDP protocols. You can also specify the Network Disk (ND) protocol, which is used by older diskless Sun workstations and the network security protocol SDNS.

By default, both UDP and ND forwarding are enabled if a helper address has been defined for an interface. The description for the **ip forward-protocol** interface configuration command in the *Cisco IOS IP and IP Routing Command Reference for Release 12.1* lists the ports that are forwarded by default if you do not specify any UDP ports.

If you do not specify any UDP ports when you configure the forwarding of UDP broadcasts, you are configuring the router to act as a BOOTP forwarding agent. BOOTP packets carry Dynamic Host Configuration Protocol (DHCP) information.

Beginning in privileged EXEC mode, follow these steps to enable forwarding UDP broadcast packets on an interface and specify the destination address:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip helper-address** *address*	Enable forwarding and specify the destination address for forwarding UDP broadcast packets, including BOOTP.
Step 4	**exit**	Return to global configuration mode.
Step 5	**ip forward-protocol** {**udp** [*port*] | **nd** | **sdns**}	Specify which protocols the router forwards when forwarding broadcast packets. • **udp**—Forward UDP datagrams. *port*: (Optional) Destination port that controls which UDP services are forwarded. • **nd**—Forward ND datagrams. • **sdns**—Forward SDNS datagrams
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show ip interface** [*interface-id*] or **show running-config**	Verify the configuration on the interface or all interfaces.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip helper-address** interface configuration command to disable the forwarding of broadcast packets to specific addresses. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

Establishing an IP Broadcast Address

The most popular IP broadcast address (and the default) is an address consisting of all ones (255.255.255.255). However, the switch can be configured to generate any form of IP broadcast address.

Beginning in privileged EXEC mode, follow these steps to set the IP broadcast address on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 3	**ip broadcast-address** *ip-address*	Enter a broadcast address different from the default, for example 128.1.255.255.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip interface** [*interface-id*]	Verify the broadcast address on the interface or all interfaces.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To restore the default IP broadcast address, use the **no ip broadcast-address** interface configuration command.

Flooding IP Broadcasts

You can allow IP broadcasts to be flooded throughout your internetwork in a controlled fashion by using the database created by the bridging STP. Using this feature also prevents loops. To support this capability, bridging must be configured on each interface that is to participate in the flooding. If bridging is not configured on an interface, it still can receive broadcasts. However, the interface never forwards broadcasts it receives, and the router never uses that interface to send broadcasts received on a different interface.

Packets that are forwarded to a single network address using the IP helper-address mechanism can be flooded. Only one copy of the packet is sent on each network segment.

To be considered for flooding, packets must meet these criteria. (Note that these are the same conditions used to consider packet forwarding using IP helper addresses.)

- The packet must be a MAC-level broadcast.
- The packet must be an IP-level broadcast.
- The packet must be a TFTP, DNS, Time, NetBIOS, ND, or BOOTP packet, or a UDP specified by the **ip forward-protocol udp** global configuration command.
- The time-to-live (TTL) value of the packet must be at least two.

A flooded UDP datagram is given the destination address specified with the **ip broadcast-address** interface configuration command on the output interface. The destination address can be set to any address. Thus, the destination address might change as the datagram propagates through the network. The source address is never changed. The TTL value is decremented.

When a flooded UDP datagram is sent out an interface (and the destination address possibly changed), the datagram is handed to the normal IP output routines and is, therefore, subject to access lists, if they are present on the output interface.

Beginning in privileged EXEC mode, follow these steps to use the bridging spanning-tree database to flood UDP datagrams:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip forward-protocol spanning-tree**	Use the bridging spanning-tree database to flood UDP datagrams.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

Use the **no ip forward-protocol spanning-tree** global configuration command to disable the flooding of IP broadcasts.

In the Catalyst 3550 switch, the majority of packets are forwarded in hardware; most packets do not go through the switch CPU. For those packets that do go to the CPU, you can speed up spanning tree-based UDP flooding by a factor of about four to five times by using turbo-flooding. This feature is supported over Ethernet interfaces configured for ARP encapsulation.

Beginning in privileged EXEC mode, follow these steps to increase spanning-tree-based flooding:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode
Step 2	**ip forward-protocol turbo-flood**	Use the spanning-tree database to speed up flooding of UDP datagrams.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To disable this feature, use the **no ip forward-protocol turbo-flood** global configuration command.

Monitoring and Maintaining IP Addressing

When the contents of a particular cache, table, or database have become or are suspected to be invalid, you can remove all its contents by using the **clear** privileged EXEC commands. Table 22-3 lists the commands for clearing contents.

Table 22-3 Commands to Clear Caches, Tables, and Databases

Command	Purpose
clear arp-cache	Clear the IP ARP cache and the fast-switching cache.
clear host {*name* | *}}	Remove one or all entries from the host name and the address cache.
clear ip route {*network* [*mask*] |*}	Remove one or more routes from the IP routing table.

You can display specific statistics, such as the contents of IP routing tables, caches, and databases; the reachability of nodes; and the routing path that packets are taking through the network. Table 22-4 lists the privileged EXEC commands for displaying IP statistics.

Table 22-4 Commands to Display Caches, Tables, and Databases

Command	Purpose
show arp	Display the entries in the ARP table.
show hosts	Display the default domain name, style of lookup service, name server hosts, and the cached list of host names and addresses.
show ip aliases	Display IP addresses mapped to TCP ports (aliases).
show ip arp	Display the IP ARP cache.
show ip interface [*interface-id*]	Display the IP status of interfaces.
show ip irdp	Display IRDP values.
show ip masks *address*	Display the masks used for network addresses and the number of subnets using each mask.
show ip redirects	Display the address of a default gateway.
show ip route [*address* [*mask*]] | [*protocol*]	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.

These are examples of outputs of commands used for displaying caches, tables, and databases.

```
Switch# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.2.3                -   0002.4b29.2e00  ARPA   GigabitEthernet0/10
Internet  172.20.136.9          178   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.250.42         207   0030.19c6.54e1  ARPA   Vlan1
Internet  120.20.30.1             -   0002.4b29.2e00  ARPA   Vlan27
Internet  172.20.139.152        159   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.139.130         24   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.141.225          5   0030.19c6.54e1  ARPA   Vlan1
Internet  172.20.135.204        227   0002.4b29.4400  ARPA   Vlan1
Internet  172.20.135.202          -   0002.4b29.2e00  ARPA   Vlan1
Internet  172.20.135.197        230   0002.4b28.ce80  ARPA   Vlan1
Internet  172.20.135.196        214   0002.4b28.ce00  ARPA   Vlan1
Internet  172.20.135.193         58   0030.19c6.54e1  ARPA   Vlan1

Switch# show hosts
Default domain is a,b,c
Name/address lookup uses static mappings

Host                    Flags      Age Type   Address(es)

Switch# show ip aliases
Address Type             IP Address      Port
Interface                10.1.2.3
Interface                120.20.30.1
Interface                172.20.135.202

Switch# show ip masks
Supernet masks  Reference count
255.252.0.0     1
0.0.0.0         1
```

```
Switch# show ip redirects
Default gateway is 172.20.135.193

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty

Switch# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
Gateway of last resort is 172.20.135.193 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 172.20.135.193
C    172.20.0.0/14 is directly connected, Vlan1

Switch# show ip route summary
IP routing table name is Default-IP-Routing-Table(0)
Route Source    Networks    Subnets     Overhead    Memory (bytes)
connected       1           0           64          144
static          1           0           64          144
rip             0           0           0           0
Total           2           0           128         288

Switch # show ip interface
Vlan1 is up, line protocol is up
  Internet address is 172.20.142.153/25
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF Fast switching turbo vector
  IP multicast fast switching is disabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are CEF
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect exclude is disabled
  BGP Policy Mapping is disabled
```

```
GigabitEthernet0/1 is up, line protocol is up
  Internet protocol processing disabled
GigabitEthernet0/2 is up, line protocol is down
  Internet protocol processing disabled
```

Enabling IP Routing

By default, the switch is in Layer 2 switching mode and IP routing is disabled. To use the Layer 3 capabilities of the switch, you must enable IP routing.

Beginning in privileged EXEC mode, follow these steps to enable IP routing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip routing**	Enable IP routing.
Step 3	**router** *ip_routing_protocol*	Specify an IP routing protocol. This step might include other commands, such as specifying the networks to route with the **network** (RIP) router configuration command. For information on specific protocols, refer to sections later in this chapter and to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip routing** global configuration command to disable routing.

This example shows how to enable IP routing using RIP as the routing protocol:

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# router rip
Switch(config-router)# network 10.0.0.0
Switch(config-router)# end
```

You can now set up parameters for the selected routing protocols as described in these sections:

- Configuring RIP, page 22-25
- Configuring IGRP, page 22-30
- Configuring OSPF, page 22-35
- Configuring EIGRP, page 22-46

You can also configure nonprotocol-specific features:

- Configuring Protocol-Independent Features, page 22-53

Configuring RIP

The Routing Information Protocol (RIP) is an interior gateway protocol (IGP) created for use in small, homogeneous networks. It is a distance-vector routing protocol that uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. The protocol is documented in RFC 1058. You can find detailed information about RIP in *IP Routing Fundamentals*, published by Cisco Press.

Using RIP, the switch sends routing information updates (advertisements) every 30 seconds. If a router does not receive an update from another router for 180 seconds or more, it marks the routes served by that router as unusable. If there is still no update after 240 seconds, the router removes all routing table entries for the non-updating router.

RIP uses hop counts to rate the value of different routes. The hop count is the number of routers that can be traversed in a route. A directly connected network has a hop count of zero; a network with a hop count of 16 is unreachable. This small range (0 to 15) makes RIP unsuitable for large networks.

If the router has a default network path, RIP advertises a route that links the router to the pseudonetwork 0.0.0.0. The 0.0.0.0 network does not exist; it is treated by RIP as a network to implement the default routing feature. The switch advertises the default network if a default was learned by RIP or if the router has a gateway of last resort and RIP is configured with a default metric. RIP sends updates to the interfaces in specified networks. If an interface's network is not specified, it is not advertised in any RIP update.

Table 22-5 shows the default RIP configuration.

Table 22-5 Default RIP Configuration

Feature	Default Setting
Auto summary	Enabled.
Default-information originate	Disabled.
Default metric	Built-in; automatic metric translations.
IP RIP authentication key-chain	No authentication. Authentication mode: clear text.
IP RIP receive version	According to the **version** router configuration command.
IP RIP send version	According to the **version** router configuration command.
IP RIP triggered	According to the **version** router configuration command.
IP split horizon	Varies with media.
Neighbor	None defined.
Network	None specified.
Offset list	Disabled.
Output delay	0 milliseconds.
Timers basic	• Update: 30 seconds. • Invalid: 180 seconds. • Hold-down: 180 seconds. • Flush: 240 seconds.

Table 22-5 Default RIP Configuration (continued)

Feature	Default Setting
Validate-update-source	Enabled.
Version	Receives RIP version 1 and version 2 packets; sends version 1 packets.

For protocol-independent features that also apply to RIP, see the "Configuring Protocol-Independent Features" section on page 22-53.

To configure RIP, you enable RIP routing for a network and optionally configure other parameters.

Beginning in privileged EXEC mode, follow these steps to enable and configure RIP:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip routing**	Enable IP routing. (Required only if IP routing is disabled.)
Step 3	**router rip**	Enable a RIP routing process, and enter router configuration mode.
Step 4	**network** *network number*	Associate a network with a RIP routing process. You can specify multiple **network** commands. RIP routing updates are sent and received through interfaces only on these networks.
Step 5	**neighbor** *ip-address*	(Optional) Define a neighboring router with which to exchange routing information. This step allows routing updates from RIP (normally a broadcast protocol) to reach nonbroadcast networks.
Step 6	**offset list** [*access-list number* | *name*] {**in** | **out**} *offset* [*type number*]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through RIP. You can limit the offset list with an access list or an interface.
Step 7	**timers basic** *update invalid holddown flush*	(Optional) Adjust routing protocol timers. Valid ranges for all timers are 0 to 4294967295 seconds. • *update*—The time (in seconds) between sending of routing updates. The default is 30 seconds. • *invalid*—The timer interval (in seconds) after which a route is declared invalid. The default is 180 seconds. • *holddown*—The time (in seconds) that must pass before a route is removed from the routing table. The default is 180 seconds. • *flush*—The amount of time (in seconds) for which routing updates are postponed. The default is 240 seconds.
Step 8	**version** {**1** | **2**}	(Optional) Configure the switch to receive and send only RIP Version 1 or RIP version 2 packets. By default, the switch receives Version 1 and 2 but sends only Version 1. You can also use the interface commands **ip rip** {**send** | **receive**} **version 1** | **2** | **1 2**} to control what versions are used for sending and receiving on interfaces.
Step 9	**no auto summary**	(Optional) Disable automatic summarization. By default, the switch summarizes subprefixes when crossing classful network boundaries. Disable summarization (RIP version 2 only) to advertise subnet and host routing information to classful network boundaries.

	Command	Purpose
Step 10	**no validate-update-source**	(Optional) Disable validation of the source IP address of incoming RIP routing updates. By default, the switch validates the source IP address of incoming RIP routing updates and discards the update if the source address is not valid. Under normal circumstances, disabling this feature is not recommended. However, if you have a router that is off-network and you want to receive its updates, you can use this command.
Step 11	**output-delay** *delay*	(Optional) Add interpacket delay for RIP updates sent. By default, packets in a multiple-packet RIP update have no delay added between packets. If you are sending packets to a lower-speed device, you can add an interpacket delay in the range of 8 to 50 milliseconds.
Step 12	**end**	Return to privileged EXEC mode.
Step 13	**show ip protocols**	Verify your entries.
Step 14	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To turn off the RIP routing process, use the **no router rip** global configuration command.

To display the parameters and current state of the active routing protocol process, use the **show ip protocols** privileged EXEC command. This is an example of output from the **show ip protocols** command, showing RIP processes:

```
Switch# show ip protocols
Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 19 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: rip
  Default version control: send version 1, receive any version
    Interface             Send  Recv  Triggered RIP  Key-chain
    Vlan1                 1     1 2
    Vlan2                 1     1 2
    GigabitEthernet0/2    1     1 2
    GigabitEthernet0/3    1     1 2                  CHAIN
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 120)
<output truncated>
```

Use the **show ip rip database** privileged EXEC command to display summary address entries in the RIP database. This example shows output with a summary address entry for route 12.11.0.0/16, with three child routes active:

```
Switch# show ip rip database
0.0.0.0/0    auto-summary
0.0.0.0/0    redistributed
    [0] via 0.0.0.0,
172.20.0.0/14    directly connected, Vlan1
```

This is an example of output of the **show ip rip database** command with a prefix and mask:

```
Switch# show ip rip database 172.19.86.0 255.255.255.0 172.19.86.0/24
    [1] via 172.19.67.38, 00:00:25, Serial0
    [1] via 172.19.70.36, 00:00:14, Serial1
```

RIP Authentication

RIP version 1 does not support authentication. If you are sending and receiving RIP Version 2 packets, you can enable RIP authentication on an interface. The key chain determines the set of keys that can be used on the interface. If a key chain is not configured, no authentication is performed, not even the default. Therefore, you must also perform the tasks in the "Managing Authentication Keys" section on page 22-63.

The switch supports two modes of authentication on interfaces for which RIP authentication is enabled: plain text and MD5. The default is plain text.

Beginning in privileged EXEC mode, follow these steps to configure RIP authentication on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 3	**ip rip authentication key-chain** *name-of-chain*	Enable RIP authentication.
Step 4	**ip rip authentication mode** [**text** | **md5**}	Configure the interface to use plain text authentication (the default) or MD5 digest authentication.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config interface** [*interface-id*]	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To restore clear text authentication, use the **no ip rip authentication mode** interface configuration command. To prevent authentication, use the **no ip rip authentication key-chain** interface configuration command.

This example shows how to verify the setting by using the **show running-config interface** privileged EXEC command.

```
Switch# show running-config interface gigabitethernet0/3
Building configuration...

Current configuration : 158 bytes
!
interface GigabitEthernet0/3
 no switchport
 ip address 10.1.3.59 255.255.255.0
 ip directed-broadcast
 ip irdp
 ip rip authentication key-chain CHAIN
end
```

Summary Addresses and Split Horizon

Routers connected to broadcast-type IP networks and using distance-vector routing protocols normally use the split-horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router on any interface from which that information originated. This feature usually optimizes communication among multiple routers, especially when links are broken.

Note In general, disabling split horizon is not recommended unless you are certain that your application requires it to properly advertise routes.

If you want to configure an interface running RIP to advertise a summarized local IP address pool on a network access server for dial-up clients, use the **ip summary-address rip** interface configuration command.

Note If split horizon is enabled, neither autosummary nor interface IP summary addresses are advertised.

Beginning in privileged EXEC mode, follow these steps to set an interface to advertise a summarized local IP address and to disable split horizon on the interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip address** *ip-address subnet-mask*	Configure the IP address and IP subnet.
Step 4	**ip summary-address rip** *ip address ip-network mask*	Configure the IP address to be summarized and the IP network mask.
Step 5	**no ip split horizon**	Disable split horizon on the interface.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show ip interface** *interface-id*	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable IP summarization, use the **no ip summary-address rip** router configuration command.

In this example, the major net is 10.0.0.0. The summary address 10.2.0.0 overrides the autosummary address of 10.0.0.0 so that 10.2.0.0 is advertised out interface Gigabit Ethernet 0.2, and 10.0.0.0 is not advertised. In the example, if the interface is still in Layer 2 mode (the default), you must enter a **no switchport** interface configuration command before entering the **ip address** interface configuration command.

Note If split horizon is enabled, neither autosummary nor interface summary addresses (those configured with the **ip summary-address rip** router configuration command) are advertised.

```
Switch(config)# router rip
Switch(config-router)# interface gi0/2
Switch(config-if)# ip address 10.1.5.1 255.255.255.0
Switch(config-if)# ip summary-address rip 10.2.0.0 255.255.0.0
Switch(config-if)# no ip split-horizon
Switch(config-if)# exit
Switch(config)# router rip
Switch(config-router)# network 10.0.0.0
Switch(config-router)# neighbor 2.2.2.2 peer-group mygroup
```

Configuring IGRP

Interior Gateway Routing Protocol (IGRP) is a dynamic, distance-vector routing, proprietary Cisco protocol for routing in an autonomous system that contains large, arbitrarily complex networks with diverse bandwidth and delay characteristics. IGRP uses a combination of user-configurable metrics, including internetwork delay, bandwidth, reliability, and load. IGRP also advertises types of routes: interior, system, and exterior, as shown in Figure 22-4.

- *Interior routes* are routes between subnets in the network attached to a router interface. If the network attached to a router is not subnetted, IGRP does not advertise interior routes.
- *System routes* are routes to networks within an autonomous system. The router derives system routes from directly connected network interfaces and system route information provided by other IGRP-speaking routers or access servers. System routes do not include subnet information.
- *Exterior routes* are routes to networks outside the autonomous system that are considered when identifying a gateway of last resort. The router chooses a gateway of last resort from the list of exterior routes that IGRP provides if it does not have a better route for a packet and the destination is not a connected network. If the autonomous system has more than one connection to an external network, different routers can choose different exterior routers as the gateway of last resort.

Figure 22-4 Interior, System, and Exterior Routes

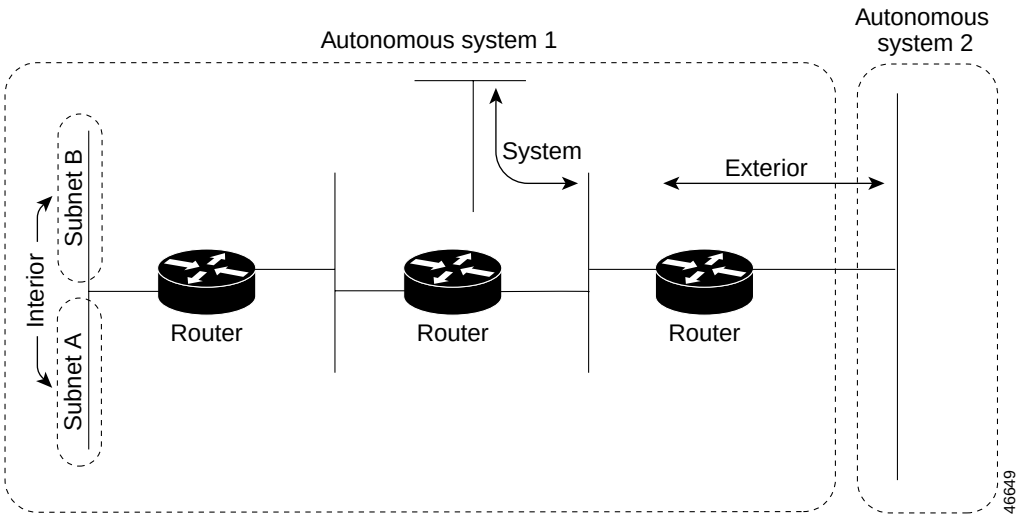

By default, a router running IGRP sends an update broadcast every 90 seconds and declares a route inaccessible if it does not receive an update from the first router in the route within three update periods (270 seconds). After seven update periods (630 seconds), the route is removed from the routing table.

Table 22-6 shows the default IGRP configuration.

Table 22-6 Default IGRP Configuration

Feature	Default Setting
IP split horizon	Varies with media.
Metric holddown	Disabled.
Metric maximum-hops	100 hops.
Neighbor	None defined.

Table 22-6 Default IGRP Configuration (continued)

Feature	Default Setting
Network	None specified.
Offset-list	Disabled.
Set metric	None set in route map.
Timers basic	Update: 90 seconds. Invalid: 270 seconds. Hold-down: 280 seconds. Flush: 630 seconds. Sleeptime: 0 milliseconds.
Traffic-share	Distributed proportionately to the ratios of the metrics.

Routers running IGRP use flash and poison-reverse updates to speed up the convergence of the routing algorithm. Flash updates are updates sent before the standard interval, notifying other routers of a metric change. Poison-reverse updates are intended to defeat larger routing loops caused by increases in routing metrics. The poison-reverse updates are sent to remove a route and place it in hold-down, which keeps new routing information from being used for a certain period of time.

Load Balancing and Traffic Distribution Control

IGRP can simultaneously use an asymmetric set of paths for a given destination. This unequal-cost load balancing allows traffic to be distributed among up to four unequal-cost paths to provide greater overall throughput and reliability.

Alternate path variance (that is, the difference in desirability between the primary and alternate paths) determines the feasibility of a potential route. An alternate route is feasible if the next router in the path is closer to the destination (has a lower metric value) than the router being used, and if the metric for the entire alternate path is within the variance. Only feasible paths are used for load balancing and are included in the routing table. These conditions limit the number of load balancing occurrences, but ensure that the dynamics of the network remain stable.

These general rules apply to IGRP unequal-cost load balancing:

- IGRP accepts up to four paths for a given destination network.
- The local best metric must be greater than the metric learned from the next router; that is, the next hop router must be closer (have a smaller metric value) to the destination than the local best metric.
- The alternative path metric must be within the specified variance of the local best metric. The multiplier times the local best metric for the destination must be greater than or equal to the metric through the next router.

If these conditions are met, the route is determined to be feasible and can be added to the routing table.

By default, the amount of variance is set to one (equal-cost load balancing). Use the **variance** router configuration command to define how much worse an alternate path can be before that path is disallowed.

If variance is configured as described in the preceding section, IGRP or Enhanced IGRP distributes traffic among multiple routes of unequal cost to the same destination. If you want faster convergence to alternate routes, but you do not want to send traffic across inferior routes in the normal case, you might prefer to have no traffic flow along routes with higher metrics. Use the **traffic-share** router configuration command to control distribution of traffic among multiple routes of unequal cost.

Note For more information and examples, refer to the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1.*

Beginning in privileged EXEC mode, follow these steps to configure IGRP. Configuring the routing process is required; other steps are optional:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router igrp** *autonomous-system*	Enable an IGRP routing process, and enter router configuration mode. The autonomous system number identifies the routes to other IGRP routers and tags routing information.
Step 3	**network** *network-number*	Associate networks with an IGRP routing process. IGRP sends updates to the interfaces in the specified networks. If an interface's network is not specified, it is not advertised in any IGRP update. It is not necessary to have a registered autonomous system number, but if you do have a registered number, we recommend that you use it to identify your process.
Step 4	**offset list** [*access-list number* | *name*] {**in** | **out**} *offset* [*type number*]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through IGRP. You can limit the offset list with an access list or an interface.
Step 5	**neighbor** *ip-address*	(Optional) Define a neighboring router with which to exchange routing information. This step allows routing updates from RIP (normally a broadcast protocol) to reach nonbroadcast network.
Step 6	**metric weights** *tos k1 k2 k3 k4 k5*	(Optional) Adjust the IGRP metric. By default, the IGRP composite metric is a 23-bit quantity that is the sum of the segment delays and the lowest segment bandwidth for a given route. • *tos*—Type of services; the default is 0. • *k1-k5*—Constants that convert a metric vector into a scalar quantity. Defaults for k1 and k3 are 1; all others are 0.
Step 7	**timers basic** *update invalid holddown flush* [*sleeptime*]	(Optional) Adjust routing protocol timers. • *update*—The time (in seconds) between sending of routing updates. The default is 90 seconds. • *invalid*—The timer interval (in seconds) after which a route is declared invalid. The default is 270 seconds. • *holddown*—The time (in seconds) during which routing information about better paths is suppressed. The default is 280 seconds. • *flush*—The time (in seconds) that must pass before a route is removed from the routing table. The default is 630 seconds. • *sleeptime*—Interval in milliseconds for postponing routing updates. The default is 0.

	Command	Purpose
Step 8	**no metric holddown**	(Optional) Disable the IGRP hold-down period. The route to a network is placed in holddown if the router learns that the network is farther away than previously known or is down. Holddown keeps new routing information from being used for a certain period of time. This can prevent routing loops caused by slow convergence. It is sometimes advantageous to disable holddown to increase the network's ability to quickly respond to topology changes; this command provides this function. Use the **metric holddown** command if other routers or access servers within the IGRP autonomous system are not configured with the **no metric holddown** command. If all routers are not configured the same way, you increase the possibility of routing loops.
Step 9	**metric maximum-hops** *hops*	(Optional) Configure the maximum network diameter. Routes with hop counts exceeding this diameter are not advertised. The default is 100 hops; the maximum is 255 hops.
Step 10	**no validate-update-source**	(Optional) Disable validation of the source IP address of incoming RIP routing updates. By default, the switch validates the source IP address of incoming RIP routing updates and discards the update if the source address is not valid.
Step 11	**variance** *multiplier*	(Optional) Define the variance associated with a particular path to enable unequal-cost load balancing if desired, balancing traffic across all feasible paths to converge to a new path if a path should fail. The multiplier can be from 1 to 128; the default is 1 (equal-cost load balancing).
Step 12	**traffic-share** {**balanced** | **min**}	(Optional) Distribute traffic by one of these methods: • **balanced**—Proportionately to the ratios of metrics • **min**—By the minimum-cost route.
Step 13	**end**	Return to privileged EXEC mode.
Step 14	**show ip protocols**	Verify your entries.
Step 15	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To shut down an IGRP routing process, use the **no router igrp** global configuration command.

This example shows how to configure a router for IGRP and assign it autonomous system 109. The **network** router configuration commands show the networks directly connected to the router.

```
Switch(config)# router igrp 109
Switch(config-router)# network 131.108.0.0
Switch(config-router)# network 192.31.7.0
```

This is an example of output from the **show ip protocols** privileged EXEC command that verifies the IGRP configuration.

```
Switch# show ip protocols
<output truncated>
Routing Protocol is "igrp 109"
  Sending updates every 90 seconds, next due in 52 seconds
  Invalid after 270 seconds, hold down 280, flushed after 630
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  IGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  IGRP maximum hopcount 100
  IGRP maximum metric variance 1
  Redistributing: igrp 109
  Maximum path: 4
  Routing for Networks:
    131.108.0.0
    183.31.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 100)
```

Split Horizon

Routers connected to broadcast-type IP networks and using distance-vector routing protocols normally use the split-horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router on any interface from which that information originated. This feature can optimize communication among multiple routers, especially when links are broken.

Note In general, we do not recommend disabling split horizon unless you are certain that your application requires it to properly advertise routes.

Beginning in privileged EXEC mode, follow these steps to disable split horizon on the interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 3	**ip address** *ip-address subnet-mask*	Configure the IP address and IP subnet.
Step 4	**no ip split-horizon**	Disable split horizon on the interface.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show ip interface** *interface-id*	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To enable the split horizon mechanism, use the **ip split-horizon** interface configuration command.

Configuring OSPF

This section briefly describes how to configure Open Shortest Path First (OSPF). For a complete description of the OSPF commands, refer to the "OSPF Commands" chapter of the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Note OSPF classifies different media into broadcast, nonbroadcast, and point-to-point networks. The Catalyst 3550 switch supports broadcast (Ethernet, Token Ring, and FDDI) and point-to-point networks (Ethernet interfaces configured as point-to-point links).

OSPF is an Interior Gateway Protocol (IGP) designed expressly for IP networks, supporting IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets. The Cisco implementation supports RFC 1253, OSPF management information base (MIB).

For protocol-independent features that include OSPF, see the "Configuring Protocol-Independent Features" section on page 22-53.

The Cisco implementation conforms to the OSPF Version 2 specifications with these key features:

- Stub areas—Definition of stub areas is supported.
- Route redistribution—Routes learned through any IP routing protocol can be redistributed into another IP routing protocol. At the intradomain level, this means that OSPF can import routes learned through IGRP and RIP. OSPF routes can also be exported into IGRP and RIP.
- Authentication—Plain text and MD5 authentication among neighboring routers within an area is supported.
- Routing interface parameter—Configurable parameters supported include interface output cost, retransmission interval, interface transmit delay, router priority, router dead and hello intervals, and authentication key.
- Virtual links—Virtual links are supported.
- Not-so-stubby-area (NSSA)—RFC 1587.

OSPF typically requires coordination among many internal routers, *area border routers* (ABRs) connected to multiple areas, and *autonomous system boundary routers* (ASBRs). The minimum configuration would use all default parameter values, no authentication, and interfaces assigned to areas. If you customize your environment, you must ensure coordinated configuration of all routers.

Table 22-7 shows the default OSPF configuration.

Table 22-7 Default OSPF Configuration

Feature	Default Setting
Interface parameters	Cost: No default cost predefined. Retransmit interval: 5 seconds. Transmit delay: 1 second. Priority: 1. Hello interval: 10 seconds. Dead interval: 4 times the hello interval. No authentication. No password specified. MD5 authentication disabled.
Area	Authentication type: 0 (no authentication). Default cost: 1. Range: Disabled. Stub: No stub area defined. NSSA: No NSSA area defined.
Auto cost	100 Mbps.
Default-information originate	Disabled. When enabled, the default metric setting is 10, and the external route type default is Type 2.
Default metric	Built-in, automatic metric translation, as appropriate for each routing protocol.
Distance OSPF	dist1 (all routes within an area): 110. dist2 (all routes from one area to another): 110. and dist3 (routes from other routing domains): 110.
OSPF database filter	Disabled. All outgoing link-state advertisements (LSAs) are flooded to the interface.
IP OSPF name lookup	Disabled.
Log adjacency changes	Enabled.
Neighbor	None specified.
Neighbor database filter	Disabled. All outgoing LSAs are flooded to the neighbor.
Network area	Disabled.
Router ID	No OSPF routing process defined.
Summary address	Disabled.
Timers LSA group pacing	240 seconds.

Table 22-7 Default OSPF Configuration (continued)

Feature	Default Setting
Timers shortest path first (spf)	spf delay: 5 seconds. spf-holdtime: 10 seconds.
Virtual link	No area ID or router ID defined. Hello interval: 10 seconds. Retransmit interval: 5 seconds. Transmit delay: 1 second. Dead interval: 40 seconds. Authentication key: no key predefined. Message-digest key (MD5): no key predefined.

Enabling OSPF requires that you create an OSPF routing process, specify the range of IP addresses to be associated with the routing process, and assign area IDs to be associated with that range.

Beginning in privileged EXEC mode, follow these steps to enable OSPF:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router ospf** *process-id*	Enable OSPF routing, and enter router configuration mode. The process ID is an internally used identification parameter that is locally assigned and can be any positive integer. Each OSPF routing process has a unique value.
Step 3	**network** *address wildcard-mask* **area** *area-id*	Define an interface on which OSPF runs and the area ID for that interface. You can use the wildcard-mask to use a single command to define one or more multiple interfaces to be associated with a specific OSPF area. The area ID can be a decimal value or an IP address.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip protocols**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To terminate an OSPF routing process, use the **no router ospf** *process-id* global configuration command.

This example shows how to configure an OSPF routing process and assign it a process number of 109:

```
Switch(config)# router ospf 109
Switch(config-router)# network 131.108.0.0 255.255.255.0 area 24
```

This is an example of output from the **show ip protocols** privileged EXEC command that verifies the OSPF process ID.

```
Switch# show ip protocols
<output truncated>
Routing Protocol is "ospf 109"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Redistributing: ospf 109
  Maximum path: 4
  Routing for Networks:
    131.108.0.0/24 Area 24
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 110)
```

OSPF Interface Parameters

You can use the **ip ospf** interface configuration commands to modify interface-specific OSPF parameters. You are not required to modify any of these parameters, but some interface parameters (hello interval, dead interval, and authentication key) must be consistent across all routers in an attached network. If you modify these parameters, be sure all routers in the network have compatible values.

Note The **ip ospf** interface configuration commands are all optional.

Beginning in privileged EXEC mode, follow these steps to modify OSPF interface parameters:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip ospf** *cost*	(Optional) Explicitly specify the cost of sending a packet on the interface.
Step 4	**ip ospf retransmit-interval** *seconds*	(Optional) Specify the number of seconds between link state advertisement transmissions. The range is 1 to 65535 seconds. The default is 5 seconds.
Step 5	**ip ospf transmit-delay** *seconds*	(Optional) Set the estimated number of seconds to wait before sending a link state update packet. The range is 1 to 65535 seconds. The default is 1 second.
Step 6	**ip ospf priority** *number*	(Optional) Set priority to help determine the OSPF designated router for a network. The range is from 0 to 255. The default is 1.
Step 7	**ip ospf hello-interval** *seconds*	(Optional) Set the number of seconds between hello packets sent on an OSPF interface. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 10 seconds.
Step 8	**ip ospf dead-interval** *seconds*	(Optional) Set the number of seconds after the last device hello packet was seen before its neighbors declare the OSPF router to be down. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 4 times the hello interval.

	Command	Purpose
Step 9	**ip ospf authentication-key** *key*	(Optional) Assign a password to be used by neighboring OSPF routers. The password can be any string of keyboard-entered characters up to 8 bytes in length. All neighboring routers on the same network must have the same password to exchange OSPF information.
Step 10	**ip ospf message digest-key** *keyid* **md5** *key*	(Optional) Enable MDS authentication. • *keyid*—An identifier from 1 to 255. • *key*—An alphanumeric password of up to 16 bytes.
Step 11	**ip ospf database-filter all out**	(Optional) Block flooding of OSPF LSA packets to the interface. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives.
Step 12	**end**	Return to privileged EXEC mode.
Step 13	**show ip ospf interface** [*interface-name*]	Display OSPF-related interface information.
Step 14	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or return to the default value.

This is an example of output from the **show ip ospf interface** privileged EXEC command:

```
Switch# show ip ospf interface
Vlan1 is up, line protocol is up
  Internet Address 172.20.135.202/14, Area 1
  Process ID 1, Router ID 172.20.135.202, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.20.135.202, Interface address 172.20.135.202
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

OSPF Area Parameters

You can optionally configure several OSPF area parameters. These parameters include authentication for password-based protection against unauthorized access to an area, stub areas, and not-so-stubby-areas (NSSAs). *Stub areas* are areas into which information on external routes is not sent. Instead, the area border router (ABR) generates a default external route into the stub area for destinations outside the autonomous system (AS). An NSSA does not flood all LSAs from the core into the area, but can import AS external routes within the area by redistribution.

Route summarization is the consolidation of advertised addresses into a single summary route to be advertised by other areas. If network numbers are contiguous, you can use the **area range** router configuration command to configure the ABR to advertise a summary route that covers all networks in the range.

Note The OSPF **area** router configuration commands are all optional.

Beginning in privileged EXEC mode, follow these steps to configure area parameters:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router ospf** *process-id*	Enable OSPF routing, and enter router configuration mode.
Step 3	**area** *area-id* **authentication**	(Optional) Allow password-based protection against unauthorized access to the identified area. The identifier can be either a decimal value or an IP address.
Step 4	**area** *area-id* **authentication message-digest**	(Optional) Enable MD5 authentication on the area.
Step 5	**area** *area-id* **stub** [**no-summary**]	(Optional) Define an area as a stub area. The **no-summary** keyword prevents an ABR from sending summary link advertisements into the stub area.
Step 6	**area** *area-id* **nssa** [**no-redistribution**] [**default-information-originate**] [**no-summary**]	(Optional) Defines an area as a not-so-stubby-area. Every router within the same area must agree that the area is NSSA. Select one of these keywords: • **no-redistribution**—Select when the router is an NSSA ABR and you want the **redistribute** command to import routes into normal areas, but not into the NSSA. • **default-information-originate**—Select on an ABR to allow importing type 7 LSAs into the NSSA. • **no-redistribution**—Select to not send summary LSAs into the NSSA.
Step 7	**area** *area-id* **range** *address mask*	(Optional) Specify an address range for which a single route is advertised. Use this command only with area border routers.
Step 8	**end**	Return to privileged EXEC mode.
Step 9	**show ip ospf** [*process-id*]	Display information about the OSPF routing process in general or for a specific process ID to verify configuration.
	show ip ospf [*process-id* [*area-id*]] **database**	Display lists of information related to the OSPF database for a specific router.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or to return to the default value.

These are examples of outputs from the **show ip ospf database** and **show ip ospf** privileged EXEC commands:

```
Switch # show ip ospf database

       OSPF Router with ID (172.20.135.202) (Process ID 1)

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.20.135.202  172.20.135.202  455         0x80000009 0x83C2   1
```

```
Switch# show ip ospf
 Routing Process "ospf 1" with ID 172.20.135.202 and Domain ID 0.0.0.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x0
 Number of opaque AS LSA 0. Checksum Sum 0x0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 External flood list length 0
    Area 1
        Number of interfaces in this area is 1
        Area has no authentication
        SPF algorithm executed 1 times
        Area ranges are
        Number of LSA 1. Checksum Sum 0x83C2
        Number of opaque link LSA 0. Checksum Sum 0x0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Other OSPF Behavior Parameters

You can optionally configure other OSPF parameters in router configuration mode.

- Route summarization: When redistributing routes from other protocols as described in the "Redistributing Routing Information" section on page 22-57, each route is advertised individually in an external LSA. To help decrease the size of the OSPF link state database, you can use the **summary-address** router configuration command to advertise a single router for all the redistributed routes included in a specified network address and mask.
- Virtual links: In OSPF, all areas must be connected to a backbone area. You can establish a virtual link in case of a backbone-continuity break by configuring two Area Border Routers as endpoints of a virtual link. Configuration information includes the identity of the other virtual endpoint (the other ABR) and the nonbackbone link that the two routers have in common (the transit area). Virtual links cannot be configured through a stub area.
- Default route: When you specifically configure redistribution of routes into an OSPF routing domain, the route automatically becomes an autonomous system boundary router (ASBR). You can force the ASBR to generate a default route into the OSPF routing domain.
- Domain Name Server (DNS) names for use in all OSPF **show** privileged EXEC command displays makes it easier to identify a router than displaying it by router ID or neighbor ID.
- Default Metrics: OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface. The metric is calculated as *ref-bw* divided by bandwidth, where *ref* is 10 by default, and bandwidth (*bw*) is determined by the **bandwidth** interface configuration command. For multiple links with high bandwidth, you can specify a larger number to differentiate the cost on those links.
- Administrative distance is a rating of the trustworthiness of a routing information source, an integer between 0 and 255, with a higher value meaning a lower trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored. OSPF uses three different administrative distances: routes within an area (interarea), routes to another area (interarea), and routes from another routing domain learned through redistribution (external). You can change any of the distance values.

- Passive interfaces: Because interfaces between two devices on an Ethernet represent only one network segment, to prevent OSPF from sending hello packets for the sending interface, you must configure the sending device to be a passive interface. Both devices can identify each other through the hello packet for the receiving interface.
- Route calculation timers: You can configure the delay time between when OSPF receives a topology change and when it starts the shortest path first (SPF) calculation and the hold time between two SPF calculations.
- Log neighbor changes: You can configure the router to send a syslog message when an OSPF neighbor state changes, providing a high-level view of changes in the router.

Beginning in privileged EXEC mode, follow these steps to configure these OSPF parameters:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router ospf** *process-id*	Enable OSPF routing, and enter router configuration mode.
Step 3	**summary-address** *address mask*	(Optional) Specify an address and IP subnet mask for redistributed routes so that only one summary route is advertised.
Step 4	**area** *area-id* **virtual-link** router-id [**hello-interval** *seconds*] [**retransmit-interval** *seconds*] [**trans**] [[**authentication-key** *key*] | **message-digest-key** *keyid* **md5** *key*]]	(Optional) Establish a virtual link and set its parameters. See the "OSPF Interface Parameters" section on page 22-38 for parameter definitions and Table 22-7 on page 22-36 for virtual link defaults.
Step 5	**default-information originate** [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-map** *map-name*]	(Optional) Force the ASBR to generate a default route into the OSPF routing domain. Parameters are all optional.
Step 6	**ip ospf name-lookup**	(Optional) Configure DNS name lookup. The default is disabled.
Step 7	**ip auto-cost reference-bandwidth** *ref-bw*	(Optional) Specify an address range for which a single route will be advertised. Use this command only with area border routers.
Step 8	**distance ospf** {[**inter-area** *dist1*] [**inter-area** *dist2*] [**external** *dist3*]}	(Optional) Change the OSPF distance values. The default distance for each type of route is 110. The range is 1 to 255.
Step 9	**passive-interface** *type number*	(Optional) Suppress the sending of hello packets through the specified interface.
Step 10	**timers spf** *spf-delay spf-holdtime*	(Optional) Configure route calculation timers. • *spf-delay*—Enter an integer from 0 to 65535. The default is 5 seconds; 0 means no delay. • *spf-holdtime*—Enter an integer from 0 to 65535. The default is 10 seconds; 0 means no delay.
Step 11	**ospf log-adj-changes**	(Optional) Send syslog message when a neighbor state changes.
Step 12	**end**	Return to privileged EXEC mode.
Step 13	**show ip ospf** [*process-id* [*area-id*]] **database**	Display lists of information related to the OSPF database for a specific router. For some of the keyword options, see to the "Monitoring OSPF" section on page 22-44.
Step 14	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Change LSA Group Pacing

The OSPF LSA group pacing feature allows the router to group OSPF LSAs and pace the refreshing, check-summing, and aging functions for more efficient router use. This feature is enabled by default with a 4-minute default pacing interval, and you will not usually need to modify this parameter. The optimum group pacing interval is inversely proportional to the number of LSAs the router is refreshing, check-summing, and aging. For example, if you have approximately 10,000 LSAs in the database, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Beginning in privileged EXEC mode, follow these steps to configure OSPF LSA pacing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router ospf** *process-id*	Enable OSPF routing, and enter router configuration mode.
Step 3	**timers lsa-group-pacing** *seconds*	Change the group pacing of LSAs.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no timers lsa-group-pacing** router configuration command.

Loopback Interface

OSPF uses the highest IP address configured on the interfaces as its router ID. If this interface is down or removed, the OSPF process must recalculate a new router ID and resend all its routing information out its interfaces. If a loopback interface is configured with an IP address, OSPF uses this IP address as its router ID, even if other interfaces have higher IP addresses. Because loopback interfaces never fail, this provides greater stability. OSPF automatically prefers a loopback interface over other interfaces, and it chooses the highest IP address among all loopback interfaces.

Beginning in privileged EXEC mode, follow these steps to configure a loopback interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface loopback 0**	Create a loopback interface, and enter interface configuration mode.
Step 3	**ip address** *address mask*	Assign an IP address to this interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip interface**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no interface loopback 0** global configuration command to disable the loopback interface.

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases.

Table 22-8 lists some of the privileged EXEC commands for displaying statistics. For more **show ip ospf database** privileged EXEC command options and for explanations of fields in the resulting display, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 22-8 Show IP OSPF Statistics Commands

Command	Purpose
show ip ospf [*process-id*]	Display general information about OSPF routing processes.
show ip ospf [*process-id*] **database** [**router**] [*link-state-id*] **show ip ospf** [*process-id*] **database** [**router**] [**self-originate**] **show ip ospf** [*process-id*] **database** [**router**] [**adv-router** [*ip-address*]] **show ip ospf** [*process-id*] **database** [**network**] [*link-state-id*] **show ip ospf** [*process-id*] **database** [**summary**] [*link-state-id*] **show ip ospf** [*process-id*] **database** [**asbr-summary**] [*link-state-id*] **show ip ospf** [*process-id*] **database** [**external**] [*link-state-id*] **show ip ospf** [*process-id area-id*] **database** [**database-summary**]	Display lists of information related to the OSPF database.
show ip ospf border-routes	Display the internal OSPF routing ABR and ASBR table entries.
show ip ospf interface [*interface-name*]	Display OSPF-related interface information.
show ip ospf neighbor [*interface-name*] [*neighbor-id*] **detail**	Display OSPF interface neighbor information.
show ip ospf virtual-links	Display OSPF-related virtual links information.

This is an example of output from the **show ip ospf** privileged EXEC command with no process ID:

```
Switch# show ip ospf
Routing Process "ospf 1" with ID 172.20.135.202 and Domain ID 0.0.0.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x0
 Number of opaque AS LSA 0. Checksum Sum 0x0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 External flood list length 0
    Area 1
        Number of interfaces in this area is 1
        Area has no authentication
        SPF algorithm executed 1 times
        Area ranges are
        Number of LSA 1. Checksum Sum 0x39E7
        Number of opaque link LSA 0. Checksum Sum 0x0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

This is an example of output from the **show ip ospf database** privileged EXEC command when no arguments or keywords are used:

```
Switch# show ip ospf database

O OSPF Router with ID (172.20.135.202) (Process ID 1)


                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
172.20.135.202  172.20.135.202  1065        0x8000002E 0x39E7   1
```

This is an example of output of the show **ip ospf interface** privileged EXEC command when Gigabit Ethernet 0/1 is specified:

```
Switch# show ip ospf interface gigabitethernet0/1

GigabitEthernet 0/1 is up, line protocol is up
Internet Address 131.119.254.202, Mask 255.255.255.0, Area 0.0.0.0
AS 201, Router ID 192.77.99.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State OTHER, Priority 1
Designated Router id 131.119.254.10, Interface address 131.119.254.10
Backup Designated router id 131.119.254.28, Interface addr 131.119.254.28
Timer intervals configured, Hello 10, Dead 60, Wait 40, Retransmit 5
Hello due in 0:00:05
Neighbor Count is 8, Adjacent neighbor count is 2
  Adjacent with neighbor 131.119.254.28  (Backup Designated Router)
  Adjacent with neighbor 131.119.254.10  (Designated Router)
```

This is an example of output from the **show ip ospf neighbor** privileged EXEC command showing a single line of summary information for each neighbor:

```
Switch# show ip ospf neighbor

   ID          Pri   State        Dead Time     Address         Interface
199.199.199.137 1    FULL/DR       0:00:31    160.89.80.37      Ethernet0
192.31.48.1     1    FULL/DROTHER  0:00:33    192.31.48.1       Fddi0
192.31.48.200   1    FULL/DROTHER  0:00:33    192.31.48.200     Fddi0
199.199.199.137 5    FULL/DR       0:00:33    192.31.48.189     Fddi0
```

This is an example of output from the **show ip ospf virtual-links** privileged EXEC command:

```
Switch# show ip ospf virtual-links

Virtual Link to router 160.89.101.2 is up
Transit area 0.0.0.1, via interface Ethernet0, Cost of using 10
Transmit Delay is 1 sec, State POINT_TO_POINT
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

Configuring EIGRP

Enhanced IGRP (EIGRP) is a Cisco proprietary enhanced version of the IGRP. Enhanced IGRP uses the same distance vector algorithm and distance information as IGRP; however, the convergence properties and the operating efficiency of Enhanced IGRP are significantly improved.

The convergence technology employs an algorithm referred to as the Diffusing Update Algorithm (DUAL), which guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations.

IP EIGRP provides increased network width. With RIP, the largest possible width of your network is 15 hops. When IGRP is enabled, the largest possible width is 224 hops. Because the EIGRP metric is large enough to support thousands of hops, the only barrier to expanding the network is the transport-layer hop counter. EIGRP increments the transport control field only when an IP packet has traversed 15 routers and the next hop to the destination was learned through EIGRP. When a RIP route is used as the next hop to the destination, the transport control field is incremented as usual.

EIGRP offers these features:

- Fast convergence.
- Incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table, minimizing the bandwidth required for EIGRP packets.
- Less CPU usage than IGRP because full update packets need not be processed each time they are received.
- Protocol-independent neighbor discovery mechanism to learn about neighboring routers.
- Variable-length subnet masks (VLSMs).
- Arbitrary route summarization.
- EIGRP scales to large networks.

Enhanced IGRP has these four basic components:

- *Neighbor discovery and recovery* is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery and recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, the Cisco IOS software can determine that a neighbor is alive and functioning. When this status is determined, the neighboring routers can exchange routing information.
- *The reliable transport protocol* is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be sent reliably, and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities (such as Ethernet), it is not necessary to send hellos reliably to all neighbors individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets (such as updates) require acknowledgment, which is shown in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. Doing so helps ensure that convergence time remains low in the presence of varying speed links.
- *The DUAL finite state machine* embodies the decision process for all route computations. It tracks all routes advertised by all neighbors. DUAL uses the distance information (known as a metric) to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. When there are no

feasible successors, but there are neighbors advertising the destination, a recomputation must occur. This is the process whereby a new successor is determined. The amount of time it takes to recompute the route affects the convergence time. Recomputation is processor-intensive; it is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL tests for feasible successors. If there are feasible successors, it uses any it finds to avoid unnecessary recomputation.

- The *protocol-dependent modules* are responsible for network layer protocol-specific tasks. An example is the IP EIGRP module, which is responsible for sending and receiving EIGRP packets that are encapsulated in IP. It is also responsible for parsing EIGRP packets and informing DUAL of the new information received. EIGRP asks DUAL to make routing decisions, but the results are stored in the IP routing table. EIGRP is also responsible for redistributing routes learned by other IP routing protocols.

Table 22-9 shows the default EIGRP configuration.

Table 22-9 Default EIGRP Configuration

Feature	Default Setting
Auto summary	Enabled. Subprefixes are summarized to the classful network boundary when crossing classful network boundaries.
Default-information	Exterior routes are accepted and default information is passed between IGRP or EIGRP processes when doing redistribution.
Default metric	Only connected routes and interface static routes can be redistributed without a default metric. The metric includes: • Bandwidth: 0 or greater kbps. • Delay (tens of microseconds): 0 or any positive number that is a multiple of 39.1 nanoseconds. • Reliability: any number between 0 and 255 (255 means 100 percent reliability). • Loading: effective bandwidth as a number between 0 and 255 (255 is 100 percent loading). • MTU: maximum transmission unit size of the route in bytes. 0 or any positive integer.
Distance	Internal distance: 90. External distance: 170.
EIGRP log-neighbor changes	Disabled. No adjacency changes logged.
IP authentication key-chain	No authentication provided.
IP authentication mode	No authentication provided.
IP bandwidth-percent	50 percent.
IP hello interval	For low-speed nonbroadcast multiaccess (NBMA) networks: 60 seconds; all other networks: 5 seconds.
IP hold-time	For low-speed NBMA networks: 180 seconds; all other networks: 15 seconds.
IP split-horizon	Enabled.
IP summary address	No summary aggregate addresses are predefined.
Metric weights	tos: 0; k1 and k3: 1; k2, k4, and k5: 0

Table 22-9 Default EIGRP Configuration (continued)

Feature	Default Setting
Network	None specified.
Offset-list	Disabled.
Router EIGRP	Disabled.
Set metric	No metric set in the route map.
Traffic-share	Distributed proportionately to the ratios of the metrics.
Variance	1 (equal-cost load balancing).

To create an EIGRP routing process, you must enable EIGRP and associate networks. EIGRP sends updates to the interfaces in the specified networks. If you do not specify an interface network, it is not advertised in any EIGRP update.

Note If you have routers on your network that are configured for IGRP, and you want to change to EIGRP, you must designate transition routers that have both IGRP and EIGRP configured. In these cases, perform Steps 1 through 3 in the next section and also see the "Configuring IGRP" section on page 22-30. You must use the same autonomous system number for routes to be automatically redistributed.

EIGRP Router Mode Commands

Beginning in privileged EXEC mode, follow these steps to configure EIGRP. Configuring the routing process is required; other steps are optional:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router eigrp** *autonomous-system*	Enable an EIGRP routing process, and enter router configuration mode. The autonomous system number identifies the routes to other EIGRP routers and is used to tag routing information.
Step 3	**network** *network-number*	Associate networks with an EIGRP routing process. EIGRP sends updates to the interfaces in the specified networks. If an interface's network is not specified, it is not advertised in any IGRP or EIGRP update.
Step 4	**eigrp log-neighbor-changes**	(Optional) Enable logging of EIGRP neighbor changes to monitor routing system stability.
Step 5	**metric weights** *tos k1 k2 k3 k4 k5*	(Optional) Adjust the EIGRP metric. Although the defaults have been carefully determined to provide excellent operation in most networks, you can adjust them. **Caution** Determining metrics is complex and is not recommended without guidance from an experienced network designer.

	Command	Purpose
Step 6	**offset list** [*access-list number* | *name*] {**in** | **out**} *offset* [*type number*]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through EIGRP. You can limit the offset list with an access list or an interface.
Step 7	**no auto-summary**	(Optional) Disable automatic summarization of subnet routes into network-level routes.
Step 8	**ip summary-address eigrp** *autonomous-system-number address mask*	(Optional) Configure a summary aggregate.
Step 9	**end**	Return to privileged EXEC mode.
Step 10	**show ip protocols**	Verify your entries.
Step 11	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

This is an example of output from the **show ip protocols** privileged EXEC command for EIGRP.

```
Switch# show ip protocols
<output truncated>
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1, igrp 1
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    172.20.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

EIGRP Interface Mode Commands

Other optional EIGRP parameters can be configured on an interface basis.

Beginning in privileged EXEC mode, follow these steps:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip bandwidth-percent eigrp** *percent*	(Optional) Configure the percentage of bandwidth that can be used by EIGRP on an interface. The default is 50 percent.
Step 4	**ip summary-address eigrp** *autonomous-system-number address mask*	(Optional) Configure a summary aggregate address for a specified interface (not usually necessary if auto-summary is enabled).

	Command	Purpose
Step 5	**ip hello-interval eigrp** *autonomous-system-number seconds*	(Optional) Change the hello time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 60 seconds for low-speed NBMA networks and 5 seconds for all other networks.
Step 6	**ip hold-time eigrp** *autonomous-system-number seconds*	(Optional) Change the hold time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 180 seconds for low-speed NBMA networks and 15 seconds for all other networks. **Caution** Do not adjust the hold time without consulting Cisco technical support.
Step 7	**no ip split-horizon eigrp** *autonomous-system-number*	(Optional) Disable split horizon to allow route information to be advertised by a router out any interface from which that information originated.
Step 8	**end**	Return to privileged EXEC mode.
Step 9	**show ip eigrp interface**	Display which interfaces EIGRP is active on and information about EIGRP relating to those interfaces.
Step 10	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

This is an example of output from the **show ip eigrp interface** privileged EXEC command:

```
Switch# show ip eigrp interface
IP-EIGRP interfaces for process 1

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Vl1            1        0/0      1000       0/10           0           0
```

Configure EIGRP Route Authentication

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol to prevent the introduction of unauthorized or false routing messages from unapproved sources.

Beginning in privileged EXEC mode, follow these steps to enable authentication:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip authentication mode eigrp** *autonomous-system* **md5**	Enable MD5 authentication in IP EIGRP packets.
Step 4	**ip authentication key-chain eigrp** *autonomous-system key-chain*	Enable authentication of IP EIGRP packets.
Step 5	**exit**	Return to global configuration mode.
Step 6	**key chain** *name-of-chain*	Identify a key chain and enter key-chain configuration mode. Match the name configured in Step 4.

	Command	Purpose
Step 7	**key** *number*	In key-chain configuration mode, identify the key number.
Step 8	**key-string** *text*	In key-chain key configuration mode, identify the key string.
Step 9	**accept-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	(Optional) Specify the time period during which the key can be received. The *start-time* and *end-time* syntax can be either *hh*:*mm*:*ss Month date year* or *hh*:*mm*:*ss date Month year*. The default is forever with the default *start-time* and the earliest acceptable date as January 1, 1993. The default *end-time* and **duration** is **infinite**.
Step 10	**send-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	(Optional) Specify the time period during which the key can be sent. The *start-time* and *end-time* syntax can be either *hh*:*mm*:*ss Month date year* or *hh*:*mm*:*ss date Month year*. The default is forever with the default *start-time* and the earliest acceptable date as January 1, 1993. The default *end-time* and **duration** is **infinite**.
Step 11	**end**	Return to privileged EXEC mode.
Step 12	**show key chain**	Display authentication key information.
Step 13	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or to return the setting to the default value.

Monitoring and Maintaining EIGRP

You can delete neighbors from the neighbor table. You can also display various EIGRP routing statistics. Table 22-10 lists the privileged EXEC commands for deleting neighbors and displaying statistics. For explanations of fields in the resulting display, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

Table 22-10 IP EIGRP Clear and Show Commands

Command	Purpose
clear ip eigrp neighbors [*if-address* | *interface*]	Delete neighbors from the neighbor table.
show ip eigrp interface [*interface*] [*as number*]	Display information about interfaces configured for EIGRP.
show ip eigrp neighbors [*type-number*]	Display EIGRP discovered neighbors.
show ip eigrp topology [*autonomous-system-number*] | [[*ip-address*] *mask*]]	Display the EIGRP topology table for a given process.
show ip eigrp traffic [*autonomous-system-number*]	Display the number of packets sent and received for all or a specified EIGRP process.

This is an example of output from the **show ip eigrp interface** privileged EXEC command:

```
Switch# show ip eigrp interface

IP EIGRP interfaces for process 109
Xmit Queue    Mean   Pacing Time   Multicast   Pending
Interface   Peers   Un/Reliable   SRTT   Un/Reliable   Flow Timer  Routes
Gi0/1         0         0/0          0      11/434          0          0
Gi0/3         1         0/0        337       0/10          0          0
Gi0/4         1         0/0         10       1/63        103          0
Gi0/5         1         0/0        330       0/16          0          0
```

This is an example of output from the **show ip eigrp neighbors** privileged EXEC command:

```
Switch# show ip eigrp neighbors

IP-EIGRP Neighbors for process 77
Address            Interface           Hold Uptime         Q        Seq    Type
                                        (secs)   (h:m:s)  Count      Num
160.89.81.28       GigabitEthernet 0/1    13      0:00:41   0         11
160.89.80.28       GigabitEthernet 0/3    14      0:02:01   0         10
160.89.80.31       GigabitEthernet 0/4    12      0:02:02   0          4
```

This is an example of output from the **show ip eigrp topology** privileged EXEC command:

```
Switch# show ip eigrp topology

IP-EIGRP Topology Table for process 77

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 160.89.90.0 255.255.255.0, 2 successors, FD is 0
          via 160.89.80.28 (46251776/46226176), GigabitEthernet0/1
          via 160.89.81.28 (46251776/46226176), GigabitEthernet0/3
          via 160.89.80.31 (46277376/46251776), GigabitEthernet0/1
P 160.89.81.0 255.255.255.0, 1 successors, FD is 307200
          via Connected, GigabitEthernet0/3
          via 160.89.81.28 (307200/281600), GigabitEthernet0/3
          via 160.89.80.28 (307200/281600), GigabitEthernet0/1
          via 160.89.80.31 (332800/307200), GigabitEthernet0/1
```

This is an example of output from the **show ip eigrp traffic** privileged EXEC command:

```
Switch# show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1
  Hellos sent/received: 19812/19821
  Updates sent/received: 1/1
  Queries sent/received: 0/0
  Replies sent/received: 0/0
  Acks sent/received: 1/0
  Input queue high water mark 1, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
```

Configuring Protocol-Independent Features

This section describes how to configure IP routing protocol-independent features. For a complete description of the IP routing protocol-independent commands in this chapter, refer to the "IP Routing Protocol-Independent Commands" chapter of the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section includes these procedures:

- Configuring Cisco Express Forwarding, page 22-53
- Configuring the Number of Equal-Cost Routing Paths, page 22-54
- Configuring Static Routes, page 22-55
- Specifying Default Routes, page 22-56
- Redistributing Routing Information, page 22-57
- Filtering Routing Information, page 22-61
- Managing Authentication Keys, page 22-63

Configuring Cisco Express Forwarding

Cisco Express Forwarding (CEF) is a Layer 3 IP switching technology used to optimize network performance. CEF implements an advanced IP look-up and forwarding algorithm to deliver maximum Layer 3 switching performance. CEF is less CPU-intensive than fast switching route caching, allowing more CPU processing power to be dedicated to packet forwarding. In the Catalyst 3550 switch, the hardware uses CEF to achieve Gigabit speed line rate IP traffic. In dynamic networks, fast switching cache entries are frequently invalidated because of routing changes, which can cause traffic to be process switched using the routing table, instead of fast switched using the route cache. CEF uses the Forwarding Information Base (FIB) lookup table to perform destination-based switching of IP packets.

The two main components in CEF are the FIB and adjacency tables.

- The FIB is similar to a routing table or information base and maintains a mirror image of the forwarding information in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table. Because the FIB contains all known routes that exist in the routing table, CEF eliminates route cache maintenance, is more efficient for switching traffic, and is not affected by traffic patterns.
- Nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.
- CEF is enabled globally by default. If for some reason it is disabled, you can re-enable it by using the **ip cef** global configuration command.

The default configuration, which is recommended, is CEF enabled on all Layer 3 interfaces. You can disable CEF on an interface by using the **no ip route-cache cef** interface configuration command; you can enable CEF on an interface by using the **ip route-cache cef** interface configuration command.

Beginning in privileged EXEC mode, follow these steps to enable CEF on an interface after it has been disabled:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	**ip route-cache cef**	Enable CEF on the interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip cef**	Display the CEF status on all interfaces.
Step 6	**show adjacency**	Display CEF adjacency table information.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable CEF on an interface, use the **no ip route-cache cef** interface configuration command.

This is an example of output from the **show ip cef** privileged EXEC command:

```
Switch# show ip cef
Prefix              Next Hop            Interface
0.0.0.0/32          receive
1.0.0.0/24          attached            GigabitEthernet0/2
1.0.0.0/32          receive
1.0.0.1/32          receive
1.0.0.55/32         1.0.0.55            GigabitEthernet0/2
1.0.0.255/32        receive
2.0.0.0/24          attached            GigabitEthernet0/3
2.0.0.0/32          receive
2.0.0.1/32          receive
2.0.0.55/32         2.0.0.55            GigabitEthernet0/3
2.0.0.255/32        receive
224.0.0.0/4         drop
224.0.0.0/24        receive
255.255.255.255/32  receive
```

This is an example of output from the **show adjacency** privileged EXEC command:

```
Switch# show adjacency
Protocol Interface                Address
IP       GigabitEthernet0/3       2.0.0.55(5)
IP       GigabitEthernet0/2       1.0.0.55(5)
```

Configuring the Number of Equal-Cost Routing Paths

When a router has two or more routes to the same network with the same metrics, these routes can be thought of as having an equal cost. The term *parallel path* is another way to refer to occurrences of equal-cost routes in a routing table. If a router has two or more equal-cost paths to a network, it can use them concurrently. Parallel paths provide redundancy in case of a circuit failure and also enables a router to load balance packets over the available paths for more efficient use of available bandwidth.

Although the router automatically learns about and configures equal-cost routes, you can control the maximum number of parallel paths supported by an IP routing protocol in its routing table.

Beginning in privileged EXEC mode, follow these steps to change the maximum number of parallel paths installed in a routing table from the default:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router** {**rip** | **ospf** | **igrp** | **eigrp**}	Enter router configuration mode.
Step 3	**maximum-paths** *maximum*	Set the maximum number of parallel paths for the protocol routing table. The range is from 1 to 8; the default is 4.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip protocols**	Verify the setting in the *Maximum path* field.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no maximum-paths** router configuration command to restore the default value.

Configuring Static Routes

Static routes are user-defined routes that cause packets moving between a source and a destination to take a specified path. Static routes can be important if the router cannot build a route to a particular destination. They are also useful for specifying a gateway of last resort to which all unroutable packets are sent.

Beginning in privileged EXEC mode, follow these steps to configure a static route:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip route** *prefix mask* {*address* | *interface*} [*distance*]	Establish a static route.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip route**	Display the current state of the routing table to verify the configuration.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip route** *prefix mask* {*address* | *interface*} global configuration command to remove a static route.

This is an example of output from the **show ip route** privileged EXEC command with a static route configured:

```
Switch# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 172.20.135.193 to network 0.0.0.0
S*   0.0.0.0/0 [1/0] via 172.20.135.193
C    172.20.0.0/14 is directly connected, Vlan1
```

The switch retains static routes until you remove them (by using the **no ip route** global configuration command). However, you can override static routes with dynamic routing information by assigning administrative distance values. Each dynamic routing protocol has a default administrative distance, as listed in Table 22-11. If you want a static route to be overridden by information from a dynamic routing protocol, set the administrative distance of the static route higher than that of the dynamic protocol.

Table 22-11 Dynamic Routing Protocol Default Administrative Distances

Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IRGP summary route	5
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
RIP	120
Unknown	225

Static routes that point to an interface are advertised through RIP, IGRP, and other dynamic routing protocols, whether or not static **redistribute** router configuration commands were specified for those routing protocols. These static routes are advertised because static routes that point to an interface are considered in the routing table to be connected and hence lose their static nature. However, if you define a static route to an interface that is not one of the networks defined in a network command, no dynamic routing protocols advertise the route unless a **redistribute** static command is specified for these protocols.

When an interface goes down, all static routes through that interface are removed from the IP routing table. When the software can no longer find a valid next hop for the address specified as the forwarding router's address in a static route, the static route is also removed from the IP routing table.

Specifying Default Routes

A router might not be able to determine the routes to all other networks. To provide complete routing capability, you can use some routers as smart routers and give the remaining routers default routes to the smart router. (Smart routers have routing table information for the entire internetwork.) These default routes can be dynamically learned or can be configured in the individual routers. Most dynamic interior routing protocols include a mechanism for causing a smart router to generate dynamic default information that is then forwarded to other routers.

Specifying a Default Network

If a router has a directly connected interface to the specified default network, the dynamic routing protocols running on that device generate or source a default route. In the case of RIP, it advertises the pseudonetwork 0.0.0.0. In the case of IGRP, the network itself is advertised and flagged as an exterior route.

A router that is generating the default for a network also might need a default of its own. One way a router can generate its own default is to specify a static route to the network 0.0.0.0 through the appropriate device.

Beginning in privileged EXEC mode, follow these steps to define a static route to a network as the static default route:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip default-network** *network number*	Specify a default network.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show ip route**	Display the selected default route in the gateway of last resort display.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no ip default-network** *network number* global configuration command to remove the route.

When default information is passed through a dynamic routing protocol, no further configuration is required. The system periodically scans its routing table to choose the optimal default network as its default route. In IGRP networks, there might be several candidate networks for the system default. Cisco routers use administrative distance and metric information to determine the default route or the gateway of last resort.

If dynamic default information is not being passed to the system, candidates for the default route are specified with the **ip default-network** global configuration command. If this network appears in the routing table from any source, it is flagged as a possible choice for the default route. If the router has no interface on the default network, but does have a path to it, the network is considered as a possible candidate, and the gateway to the best default path becomes the gateway of last resort.

Redistributing Routing Information

The switch can run multiple routing protocols simultaneously, and it can redistribute information from one routing protocol to another. For example, you can instruct the switch to readvertise IGRP-derived routes by using RIP or to readvertise static routes by using IGRP. Redistributing information from one routing protocol to another applies to all supported IP-based routing protocols.

You can also conditionally control the redistribution of routes between routing domains by defining a method known as route maps between the two domains. Although redistribution is a protocol-independent feature, some of the **match** and **set** route-map configuration commands are specific to a particular protocol.

One or more **match** commands and one or more **set** commands follow a **route-map** command. If there are no **match** commands, everything matches. If there are no **set** commands, nothing is done, other than the match. Therefore, you need at least one **match** or **set** command.

Note Although each of Steps 3 through 14 in the following section is optional, you must enter at least one **match** route-map configuration command and one **set** route-map configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a route map for redistribution:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**route-map** *map-tag* [**permit** | **deny**] [*sequence number*]	Define any route maps used to control redistribution and enter route-map configuration mode. *map-tag*—A meaningful name for the route map. The **redistribute** router configuration command uses this name to reference this route map. Multiple route maps might share the same map tag name. (Optional) If **permit** is specified and the match criteria are met for this route map, the route is redistributed as controlled by the set actions. If **deny** is specified, the route is not redistributed. *sequence number* (Optional)— Number that indicates the position a new route map is to have in the list of route maps already configured with the same name.
Step 3	**match ip address** {*access-list-number* | *access-list-name*} [*...access-list-number* | *...access-list-name*]	Match a standard access list by specifying the name or number. It can be an integer from 1 to 199.
Step 4	**match metric** *metric-value*	Match the specified route metric. The *metric-valu*e can be an IGRP five-part metric with a specified value from 0 to 4294967295.
Step 5	**match ip next-hop** {*access-list-number* | *access-list-name*} [*...access-list-number* | *...access-list-name*]	Match a next-hop router address passed by one of the access lists specified (numbered from 1 to 199).
Step 6	**match tag** *tag value* [*...tag-value*]	Match the specified tag value in a list of one or more route tag values. Each can be an integer from 0 to 4294967295.
Step 7	**match interface** *type number* [*...type number*]	Match the specified next hop route out one of the specified interfaces.
Step 8	**match ip route-source** {*access-list-number* | *access-list-name*} [*...access-list-number* | *...access-list-name*]	Match the address specified by the specified advertised access lists.
Step 9	**match route-type** {**local** | **internal** | **external** [**type-1** | **type-2**] | **level-1** | **level-2**}	Match the specified **route-type**: • **local**—Locally generated BGP routes. • **internal**—OSPF intra-area and interarea routes or EIGRP internal routes. • **external**—OSPF external routes (Type 1 or Type 2) or EIGRP external routes. • **level-1 | -2**—IS-IS Level 1 or Level 2 routes.
Step 10	**set next-hop** *next-hop*	Specify the address of the next hop.
Step 11	**set level** {**level-1** | **level-2** | **level-1-2** | **stub-area** | **backbone**}	Set the level for routes that are advertised into the specified area of the routing domain. The **stub-area** and **backbone** are OSPF NSSA and backbone areas.

	Command	Purpose
Step 12	**set metric** *metric value*	Set the metric value to give the redistributed routes (for any protocol except IGRP or EIGRP). The *metric value* is an integer from -294967295 to 294967295.
Step 13	**set metric** *bandwidth delay reliability loading mtu*	Set the metric value to give the redistributed routes (for IGRP or EIGRP only): • *bandwidth*—Metric value or IGRP bandwidth of the route in kilobits per second in the range 0 to 4294967295 • *delay*—Route delay in tens of microseconds in the range 0 to 4294967295. • *reliability*—Likelihood of successful packet transmission expressed as a number between 0 and 255, where 255 means 100 percent reliability and 0 means no reliability. • *loading*— Effective bandwidth of the route expressed as a number from 0 to 255 (255 is 100 percent loading). • *mtu*—Minimum maximum transmission unit (MTU) size of the route in bytes in the range 0 to 4294967295.
Step 14	**set metric-type** {**internal** | **external** | **type-1** | **type-2**}	Set the metric type to give redistributed routes.
Step 15	**end**	Return to privileged EXEC mode.
Step 16	**show route-map**	Display all route maps configured or only the one specified to verify configuration.
Step 17	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To delete an entry, use the **no route-map** *map tag* global configuration command or the **no match** or **no set** route-map configuration commands.

You can distribute routes from one routing domain into another and control route distribution.

Beginning in privileged EXEC mode, follow these steps to control route redistribution. Note that the keywords are the same as defined in the previous procedure.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router** {**rip** | **ospf** | **igrp** | **eigrp**}	Enter router configuration mode.
Step 3	**redistribute** *protocol* [*process-id*] {**level-1** | **level-1-2** | **level-2**} [**metric** *metric-value*] [**metric-type** *type-value*] [**match internal** | **external** *type-value*] [**tag** *tag-value*] [**route-map** *map-tag*] [**weight** *weight*] [**subnets**]	Redistribute routes from one routing protocol to another routing protocol.
Step 4	**default-metric** *number*	Cause the current routing protocol to use the same metric value for all redistributed routes (RIP and OSPF).
Step 5	**default-metric** *bandwidth delay reliability loading mtu*	Cause the IGRP or EIGRP routing protocol to use the same metric value for all non-IGRP redistributed routes.
Step 6	**no default-information** {**in** | **out**}	Disable the redistribution of default information between IGRP processes, which is enabled by default.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show route-map**	Display all route maps configured or only the one specified to verify configuration.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable redistribution, use the **no** form of the commands.

The metrics of one routing protocol do not necessarily translate into the metrics of another. For example, the RIP metric is a hop count, and the IGRP metric is a combination of five qualities. In these situations, an artificial metric is assigned to the redistributed route. Uncontrolled exchanging of routing information between different routing protocols can create routing loops and seriously degrade network operation.

If you have not defined a default redistribution metric that replaces metric conversion, some automatic metric translations occur between routing protocols:

- RIP can automatically redistribute static routes. It assigns static routes a metric of 1 (directly connected).
- IGRP can automatically redistribute static routes and information from other IGRP-routed autonomous systems. IGRP assigns static routes a metric that identifies them as directly connected. It does not change the metrics of routes derived from IGRP updates from other autonomous systems.
- Any protocol can redistribute other routing protocols if a default mode is in effect.

Filtering Routing Information

You can filter routing protocol information by performing the tasks described in this section.

Note When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

Setting Passive Interfaces

To prevent other routers on a local network from dynamically learning about routes, you can use the **passive-interface** router configuration command to keep routing update messages from being sent through a router interface. When you use this command in the OSPF protocol, the interface address you specify as passive appears as a stub network in the OSPF domain. OSPF routing information is neither sent nor received through the specified router interface.

In networks with many interfaces, to avoid having to manually set them as passive, you can set all interfaces to be passive by default by using the **passive-interface default** router configuration command and manually setting interfaces where adjacencies are desired.

Beginning in privileged EXEC mode, follow these steps to configure passive interfaces:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router** {**rip** | **ospf** | **igrp** | **eigrp**}	Enter router configuration mode.
Step 3	**passive-interface** *interface-id*	Suppress sending routing updates through the specified Layer 3 interface.
Step 4	**passive-interface default**	(Optional) Set all interfaces as passive by default.
Step 5	**no passive-interface** *interface type*	(Optional) Activate only those interfaces that need to have adjacencies sent.
Step 6	**network** *network-address*	(Optional) Specify the list of networks for the routing process. The *network-address* is an IP address.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use a network monitoring privileged EXEC command such as **show ip ospf interface** to verify the interfaces that you enabled as passive, or use the **show ip interface** privileged EXEC command to verify the interfaces that you enabled as active.

To re-enable the sending of routing updates, use the **no passive-interface** *interface-id* router configuration command. The **default** keyword sets all interfaces as passive by default. You can then configure individual interfaces where you want adjacencies by using the **no passive-interface** router configuration command. The **default** keyword is useful in Internet service provider and large enterprise networks where many of the distribution routers have more than 200 interfaces.

Controlling Advertising and Processing in Routing Updates

You can use the **distribute-list** router configuration command with access control lists to suppress routes from being advertised in routing updates and to prevent other routers from learning one or more routes. When used in OSPF, this feature applies to only external routes, and you cannot specify an interface name.

You can also use a **distribute-list** router configuration command to avoid processing certain routes listed in incoming updates. (This feature does not apply to OSPF.)

Beginning in privileged EXEC mode, follow these steps to control the advertising or processing of routing updates:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router** {**rip** | **ospf** | **igrp** | **eigrp**}	Enter router configuration mode.
Step 3	**distribute-list** {*access-list-number* | *access-list-name*} **out** [*interface-name* | *routing process* | *autonomous-system-number*]	Permit or deny routes from being advertised in routing updates, depending upon the action listed in the access list.
Step 4	**distribute-list** {*access-list-number* | *access-list-name*} **in** [*type-number*]	Suppress processing in routes listed in updates.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no distribute-list in** router configuration command to change or cancel a filter. To cancel suppression of network advertisements in updates, use the **no distribute-list out** router configuration command.

Filtering Sources of Routing Information

Because some routing information might be more accurate than others, you can use filtering to prioritize information coming from different sources. An *administrative distance* is a rating of the trustworthiness of a routing information source, such as a router or group of routers. In a large network, some routing protocols can be more reliable than others. By specifying administrative distance values, you enable the router to intelligently discriminate between sources of routing information. The router always picks the route whose routing protocol has the lowest administrative distance. Table 22-11 on page 22-56 shows the default administrative distances for various routing information sources.

Because each network has its own requirements, there are no general guidelines for assigning administrative distances.

Beginning in privileged EXEC mode, follow these steps to filter sources of routing information:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**router** {**rip** | **ospf** | **igrp** | **eigrp**}	Enter router configuration mode.

	Command	Purpose
Step 3	**distance** *weight* {*ip-address* {*ip-address mask*}} [*ip access list*]	Define an administrative distance. *weight*—The administrative distance as an integer from 10 to 255. Used alone, *weight* specifies a default administrative distance that is used when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. (Optional) *ip access list*—An IP standard or extended access list to be applied to incoming routing updates.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip protocols**	Display the default administrative distance for a specified routing process.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a distance definition, use the **no distance** router configuration command.

Managing Authentication Keys

Key management is a method of controlling authentication keys used by routing protocols. Not all protocols can use key management. Authentication keys are available for EIGRP and RIP Version 2.

Before you manage authentication keys, you must enable authentication. See the appropriate protocol section to see how to enable authentication for that protocol. To manage authentication keys, define a key chain, identify the keys that belong to the key chain, and specify how long each key is valid. Each key has its own key identifier (specified with the **key** *number* key chain configuration command), which is stored locally. The combination of the key identifier and the interface associated with the message uniquely identifies the authentication algorithm and Message Digest 5 (MD5) authentication key in use.

You can configure multiple keys with life times. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest, and uses the first valid key it encounters. The lifetimes allow for overlap during key changes. Note that the router must know these lifetimes.

Beginning in privileged EXEC mode, follow these steps to manage authentication keys:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**key chain** *name-of-chain*	Identify a key chain, and enter key chain configuration mode.
Step 3	**key** *number*	Identify the key number. The range is 0 to 2147483647.
Step 4	**key-string** *text*	Identify the key string. The string can contain from 1 to 80 uppercase and lowercase alphanumeric characters, but the first character cannot be a number.

	Command	Purpose
Step 5	**accept-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	(Optional) Specify the time period during which the key can be received. The *start-time* and *end-time* syntax can be either *hh*:*mm*:*ss Month date year* or *hh*:*mm*:*ss date Month year*. The default is forever with the default *start-time* and the earliest acceptable date as January 1, 1993. The default *end-time* and **duration** is **infinite**.
Step 6	**send-lifetime** *start-time* {**infinite** | *end-time* | **duration** *seconds*}	(Optional) Specify the time period during which the key can be sent. The *start-time* and *end-time* syntax can be either *hh*:*mm*:*ss Month date year* or *hh*:*mm*:*ss date Month year*. The default is forever with the default *start-time* and the earliest acceptable date as January 1, 1993. The default *end-time* and **duration** is infinite.
Step 7	**end**	Return to privileged EXEC mode.
Step 8	**show key chain**	Display authentication key information.
Step 9	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the key chain, use the **no key chain** *name-of-chain* global configuration command.

Monitoring and Maintaining the IP Network

You can remove all contents of a particular cache, table, or database. You can also display specific statistics. Use the privileged EXEC commands in Table 22-12 to clear routes or display status:

Table 22-12 Commands to Clear IP Routes or Display Route Status

Command	Purpose
clear ip route {*network* [*mask* | *]}	Clear one or more routes from the IP routing table.
show ip protocols	Display the parameters and state of the active routing protocol process.
show ip route [*address* [*mask*] [**longer-prefixes**]] | [*protocol* [*process-id*]]	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.
show ip route supernets-only	Display supernets.
show ip cache	Display the routing table used to switch IP traffic.
show route-map [*map-name*]	Display all route maps configured or only the one specified.

This is an example of output from the **show ip route** privileged EXEC command when entered without an address:

```
Switch# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 172.20.135.193 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 172.20.135.193
C    172.20.0.0/14 is directly connected, Vlan1
```

This is an example of output from the **show ip protocols** privileged EXEC command, showing IGRP processes:

```
Switch# show ip protocols

Routing Protocol is "igrp 1"
  Sending updates every 90 seconds, next due in 65 seconds
  Invalid after 270 seconds, hold down 280, flushed after 630
  Outgoing update filter list for all interfaces is
  Incoming update filter list for all interfaces is
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  IGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  IGRP maximum hopcount 100
  IGRP maximum metric variance 1
  Redistributing: eigrp 1, igrp 1
  Maximum path: 4
  Routing for Networks:
    172.20.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 100)
```

This is an example of output from the **show ip route summary** privileged EXEC command:

```
Switch# show ip route summary
IP routing table name is Default-IP-Routing-Table(0)
Route Source    Networks    Subnets     Overhead    Memory (bytes)
connected       1           0           64          144
static          1           0           64          144
rip             0           0           0           0
eigrp 1         0           0           0           0
igrp 1          0           0           0           0
ospf 1          0           0           0           0
  Intra-area: 0 Inter-area: 0 External-1: 0 External-2: 0
  NSSA External-1: 0 NSSA External-2: 0
Total           2           0           128         288
```

This is an example of output from the **show ip route supernets-only** privileged EXEC command. This display shows supernets only; it does not show subnets.

```
Switch# show ip route supernets-only

Codes: I - IGRP derived, R - RIP derived, O - OSPF derived
       C - connected, S - static, E - EGP derived, B - BGP derived
       i - IS-IS derived, D - EIGRP derived
       * - candidate default route, IA - OSPF inter area route
       E1 - OSPF external type 1 route, E2 - OSPF external type 2 route
       L1 - IS-IS level-1 route, L2 - IS-IS level-2 route
       EX - EIGRP external route

Gateway of last resort is not set

B    198.92.0.0 (mask is 255.255.0.0) [20/0] via 198.92.72.30, 0:00:50
B    192.0.0.0 (mask is 255.0.0.0) [20/0] via 198.92.72.24, 0:02:50
```

This is an example of output from the **show route-map** privileged EXEC command:

```
Switch# show route-map

route-map abc, permit, sequence 10
  Match clauses:
    tag 1 2
  Set clauses:
    metric 5
route-map xyz, permit, sequence 20
  Match clauses:
    tag 3 4
  Set clauses:
    metric 6
```

CHAPTER 23

Configuring HSRP

This chapter describes how to use Hot Standby Router Protocol (HSRP) to provide routing redundancy for routing IP traffic without being dependent on the availability of any single router. To use this feature, you must have the enhanced multilayer software image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Note You can also use a version of HSRP in Layer 2 mode to configure a redundant command switch to take over cluster management if the cluster command switch fails. For more information about clustering, see Chapter 5, "Clustering Switches."

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* and the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding HSRP, page 23-1
- Configuring HSRP, page 23-3
- Displaying HSRP Configurations, page 23-10

Understanding HSRP

HSRP is Cisco's standard method of providing high network availability by providing first-hop redundancy for IP hosts on an IEEE 802 LAN configured with a default gateway IP address. HSRP routes IP traffic without relying on the availability of any single router. It enables a set of router interfaces to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. When HSRP is configured on a network or segment, it provides a virtual Media Access Control (MAC) address and an IP address that is shared among a group of configured routers. HSRP allows two or more HSRP-configured routers to use the MAC address and IP network address of a virtual router. The virtual router does not exist; it represents the common target for routers that are configured to provide backup to each other. One of the routers is selected to be the active router and another to be the standby router, which assumes control of the group MAC address and IP address should the designated active router fail.

Note Routers in an HSRP group can be any router interface that supports HSRP, including Catalyst 3550 routed ports and switch virtual interfaces (SVIs).

HSRP provides high network availability by providing redundancy for IP traffic from hosts on networks. In a group of router interfaces, the active router is the router of choice for routing packets; the standby router is the router that takes over the routing duties when an active router fails or when preset conditions are met.

HSRP is useful for hosts that do not support a router discovery protocol and cannot switch to a new router when their selected router reloads or loses power. When HSRP is configured on a network segment, it provides a virtual MAC address and an IP address that is shared among router interfaces in a group of router interfaces running HSRP. The router selected by the protocol to be the active router receives and routes packets destined for the group's MAC address. For *n* routers running HSRP, there are *n +1* IP and MAC addresses assigned.

HSRP detects when the designated active router fails, and a selected standby router assumes control of the Hot Standby group's MAC and IP addresses. A new standby router is also selected at that time. Devices running HSRP send and receive multicast UDP-based hello packets to detect router failure and to designate active and standby routers. When HSRP is configured on an interface, Internet Control Message Protocol (ICMP) redirect messages are disabled by default for the interface.

You can configure multiple Hot Standby groups among Catalyst 3550 switches that are operating in Layer 3 to make more use of the redundant routers. To do so, specify a group number for each Hot Standby command group you configure for an interface. For example, you might configure an interface on switch 1 as an active router and one on switch 2 as a standby router and also configure another interface on switch 2 as an active router with another interface on switch 1 as its standby router.

Figure 23-1 shows a segment of a network configured for HSRP. Each router is configured with the MAC address and IP network address of the virtual router. Instead of configuring hosts on the network with the IP address of Router A, you configure them with the IP address of the virtual router as their default router. When Host C sends packets to Host B, it sends them to the MAC address of the virtual router. If for any reason, Router A stops transferring packets, Router B responds to the virtual IP address and virtual MAC address and becomes the active router, assuming the active router duties. Host C continues to use the IP address of the virtual router to address packets destined for Host B, which Router B now receives and sends to Host B. Until Router A resumes operation, HSRP allows Router B to provide uninterrupted service to users on Host C's segment that need to communicate with users on Host B's segment and also continues to perform its normal function of handling packets between the Host A segment and Host B.

Figure 23-1 Typical HSRP Configuration

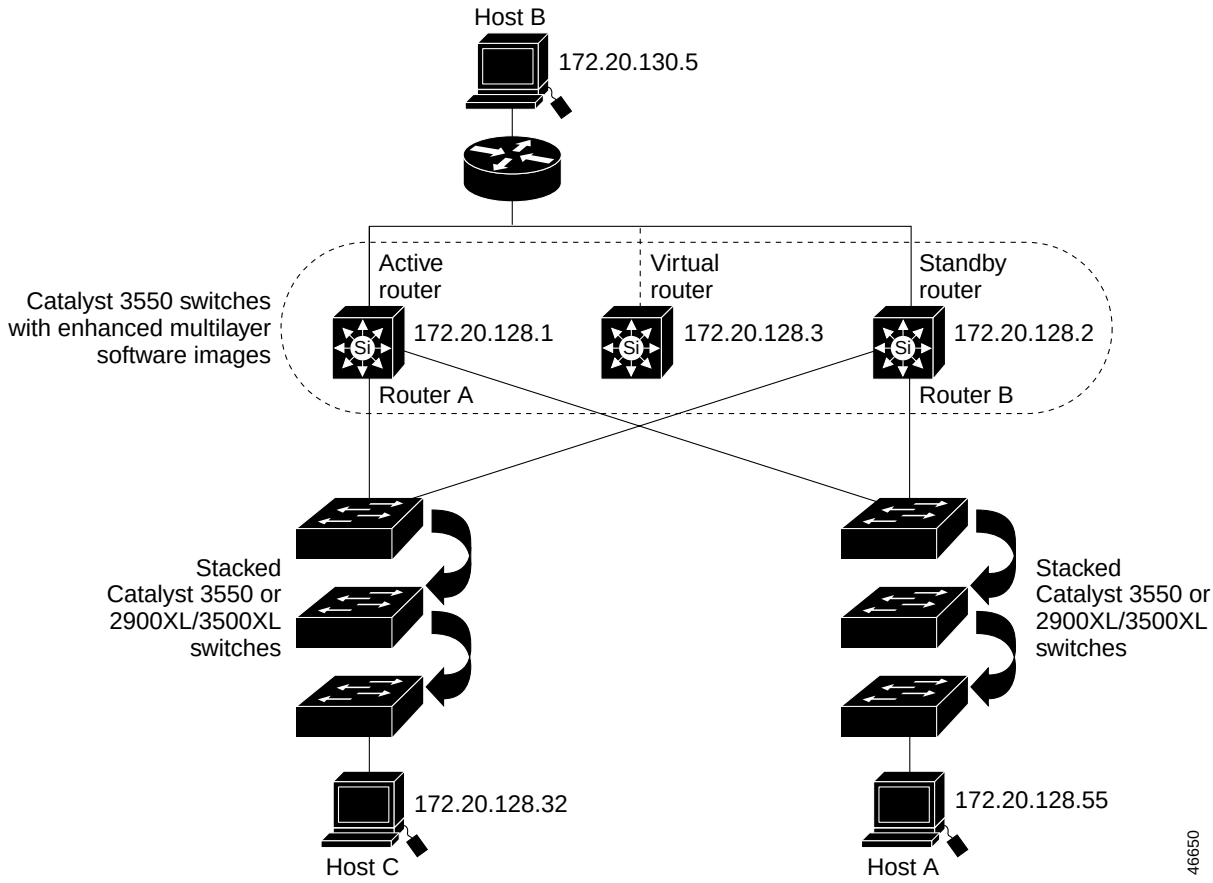

Configuring HSRP

These sections include HSRP configuration information:

- Default HSRP Configuration, page 23-4
- Enabling HSRP, page 23-4
- Configuring HSRP Group Attributes, page 23-6
- Configuring HSRP Groups and Clustering, page 23-9

Note If HSRP is enabled, the switch can recognize 16 additional MAC addresses, each associated with a set of VLANs or routing interfaces.

In the following procedures, the specified interface must be one of these Layer 3 interfaces:

- Routed port: a physical port configured as a Layer 3 port by entering the **no switchport** interface configuration command.
- SVI: a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface.

- Etherchannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel** *port-channel-number* global configuration command and binding the Ethernet interface into the channel group. For more information, see the "Configuring Layer 3 EtherChannels" section on page 21-11.

All Layer 3 interfaces must have IP addresses assigned to them. See the "Configuring Layer 3 Interfaces" section on page 8-22.

Default HSRP Configuration

Table 23-1 shows the default HSRP configuration.

Table 23-1 Default HSRP Configuration

Feature	Default Setting
HSRP groups	None configured
Standby group number	0
Standby MAC address	Well-known MAC address
Standby priority	100
Standby delay	0 (no delay)
Standby track interface priority	10
Standby hello time	3 seconds
Standby holdtime	10 seconds

Enabling HSRP

The **standby ip** interface configuration command activates HSRP on the configured interface. If an IP address is specified, that address is used as the designated address for the Hot Standby group. If no IP address is specified, the address is learned through the standby function. You must configure at least one routing port on the cable with the designated address. Configuring an IP address always overrides another designated address currently in use.

When the **standby ip** command is enabled on an interface and proxy ARP is enabled, if the interface's Hot Standby state is active, proxy ARP requests are answered using the Hot Standby group MAC address. If the interface is in a different state, proxy ARP responses are suppressed.

Note We recommend that you do not configure a virtual MAC address on the switch interfaces, but instead use the default well-known virtual MAC address. If it is necessary to configure a virtual MAC address, for HSRP to operate correctly, the first five bytes must be the same for all MAC addresses used for HSRP. The last byte represents the group number.

Beginning in privileged EXEC mode, follow these steps to create or enable HSRP on a Layer 3 interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the Layer 3 interface on which you want to enable HSRP.
Step 3	**standby** [*group-number*] **ip** [*ip-address* [**secondary**]]	Create (or enable) the HSRP group using its number and virtual IP address. • (Optional) *group-number*—The group number on the interface for which HSRP is being enabled. The range is 0 to 255; the default is 0. If there is only one HSRP group, you do not need to enter a group number. • (Optional on all but one interface) *ip-address*—The virtual IP address of the hot standby router interface. You must enter the virtual IP address for at least one of the interfaces; it can be learned on the other interfaces. • (Optional) **secondary**—The IP address is a secondary hot standby router interface. If neither router is designated as a secondary or standby router and no priorities are set, the primary IP addresses are compared and the higher IP address is the active router, with the next highest as the standby router.
Step 4	**standby** [*group-number*] **mac-address** *mac-address*	(Optional) Specify the virtual MAC address for the virtual router. We recommend that you do not configure this parameter, but instead use the default well-known MAC address. If it is necessary to configure a virtual MAC address, the first five bytes in all MAC addresses for the HSRP should be the same. The last byte represents the group number.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show standby** [*interface-id* [*group*]]	Verify the configuration.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no standby** [*group-number*] **ip** [*ip-address*] interface configuration command to disable HSRP.

Use the **no standby** [*group-number*] **mac-address** *mac-address* interface configuration command to return to the default virtual MAC address.

This example shows how to activate HSRP for group 1 on Gigabit Ethernet interface 0/1. The IP address used by the hot standby group is learned by using HSRP.

Note This procedure is the minimum number of steps required to enable HSRP.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# standby 1 ip
Switch(config-if)# end
Switch# show standby
```

Configuring HSRP Group Attributes

Although HSRP can run with no other configuration required, you can configure attributes for the HSRP group, including authentication, priority, preemption and preemption delay, timers, or MAC address.

Configuring HSRP Priority

The **standby priority**, **standby preempt**, and **standby track** interface configuration commands are all used to set characteristics for determining active and standby routers and behavior regarding when a new active router takes over. When configuring priority, follow these guidelines:

- Assigning priority helps select the active and standby routers. If preemption is enabled, the router with the highest priority becomes the designated active router. If priorities are equal, the primary IP addresses are compared, and the higher IP address has priority.
- The highest number (1 to 255) represents the highest priority (most likely to become the active router).
- When setting the priority, preempt, or both, you must specify at least one keyword (**priority**, **preempt**, or both).
- The priority of the device can change dynamically if an interface is configured with the **standby track** command and another interface on the router goes down.
- The **standby track** interface configuration command ties the router hot standby priority to the availability of its interfaces and is useful for tracking interfaces that are not configured for HSRP. When a tracked interface fails, the hot standby priority on the device on which tracking has been configured decreases by 10. If an interface is not tracked, its state changes do not affect the hot standby priority of the configured device. For each interface configured for hot standby, you can configure a separate list of interfaces to be tracked.
- The **standby track** *interface-priority* interface configuration command specifies how much to decrement the hot standby priority when a tracked interface goes down. When the interface comes back up, the priority is incremented by the same amount.
- When multiple tracked interfaces are down and *interface-priority* values have been configured, the configured priority decrements are cumulative. If tracked interfaces that were not configured with priority values fail, the default decrement is 10, and it is noncumulative.
- When routing is first enabled for the interface, it does not have a complete routing table. If it is configured to preempt, it becomes the active router, even though it is unable to provide adequate routing services. To solve this problem, configure a delay time to allow the router to update its routing table.

Beginning in privileged EXEC mode, use one or more of these steps to configure HSRP priority characteristics on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the HSRP interface on which you want to set priority.
Step 3	**standby** [*group-number*] **priority** *priority* [**preempt** [**delay** *delay*]]	Set a **priority** value used in choosing the active router. The range is 1 to 255; the default priority is 100. The highest number represents the highest priority. • (Optional) *group-number*—The group number to which the command applies. • (Optional) **preempt—**Select so that when the local router has a higher priority than the active router, it assumes control as the active router. • (Optional) **delay**—Set to cause the local router to postpone taking over the active role for the shown number of seconds. The range is 0 to 36000 (1 hour); the default is 0 (no delay before taking over). Use the **no** form of the command to restore the default values.
Step 4	**standby** [*group-number*] [**priority** *priority*] **preempt** [**delay** *delay*]	Configure the router to **preempt**, which means that when the local router has a higher priority than the active router, it assumes control as the active router. • (Optional) *group-number*—The group number to which the command applies. • (Optional) **priority**—Enter to set or change the group priority. The range is 1 to 255; the default is 100. • (Optional) **delay**—Set to cause the local router to postpone taking over the active role for the number of seconds shown. The range is 0 to 36000 (1 hour); the default is 0 (no delay before taking over). Use the **no** form of the command to restore the default values.
Step 5	**standby** [*group-number*] **track** *type number* [*interface-priority*]	Configure an interface to track other interfaces so that if one of the other interfaces goes down, the device's Hot Standby priority is lowered. • (Optional) *group-number*—The group number to which the command applies. • *type*—Enter the interface type (combined with interface number) that is tracked. • *number*—Enter the interface number (combined with interface type) that is tracked. • (Optional) *interface-priority*—Enter the amount by which the hot standby priority for the router is decremented or incremented when the interface goes down or comes back up. The default value is 10.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify the configuration of the standby groups.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no standby** [*group-number*] **priority** *priority* [**preempt** [**delay** *delay*]] and **no standby** [*group-number*] [**priority** *priority*] **preempt** [**delay** *delay*] interface configuration commands to restore default priority, preempt, and delay values.

Use the **no standby** [*group-number*] **track** *type number* [*interface-priority*] interface configuration command to remove the tracking.

This activates Gigabit Ethernet interface 0/1, sets an IP address and a priority of 120 (higher than the default value), and waits for 300 seconds (5 minutes) before attempting to become the active router:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# standby ip 172.19.108.254
Switch(config-if)# standby priority 120 preempt delay 300
Switch(config-if)# end
Switch#
```

Configuring HSRP Authentication and Timers

You can optionally configure an HSRP authentication string or change the hello-time interval and holdtime.

When configuring these attributes, follow these guidelines:

- The authentication string is sent unencrypted in all HSRP messages. You must configure the same authentication string on all routers and access servers on a cable to ensure interoperation. Authentication mismatch prevents a device from learning the designated Hot Standby IP address and timer values from other routers configured with HSRP.
- Routers or access servers on which standby timer values are not configured can learn timer values from the active or standby router. The timers configured on an active router always override any other timer settings.
- All routers in a Hot Standby group should use the same timer values. Normally, the *holdtime* is greater than or equal to 3 times the *hellotime*.

Beginning in privileged EXEC mode, use one or more of these steps to configure HSRP priority characteristics on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and enter the HSRP interface on which you want to set authentication.
Step 3	**standby** [*group-number*] **authentication** *string*	(Optional) **authentication** *string*—Enter a string to be carried in all HSRP messages. The authentication string can be up to eight characters in length; the default string is **cisco.** (Optional) *group-number*—The group number to which the command applies.

	Command	Purpose
Step 4	**standby** [*group-number*] **timers** *hellotime holdtime*	(Optional) Configure the time between hello packets and the time before other routers declare the active router to be down. • *group-number*—The group number to which the command applies. • *hellotime*—The hello interval in seconds. The range is from 1 to 255; the default is 3 seconds. • *holdtime*—The time in seconds before the active or standby router is declared to be down. The range is from 1 to 255; the default is 10 seconds.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify the configuration of the standby groups.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Use the **no standby** [*group-number*] **authentication** *string* interface configuration command to delete an authentication string. Use the **no standby** [*group-number*] **timers** *hellotime holdtime* interface configuration command to restore timers to their default values.

This example shows how to configure *word* as the authentication string required to allow Hot Standby routers in group 1 to interoperate:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# standby 1 authentication word
Switch(config-if)# end
Switch#
```

This example shows how to set the timers on standby group 1 with the time between hello packets at 5 seconds and the time after which a router is considered down to be 15 seconds:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# standby 1 ip
Switch(config-if)# standby 1 timers 5 15
Switch(config-if)# end
Switch#
```

Configuring HSRP Groups and Clustering

When a device is participating in an HSRP standby routing and clustering is enabled, you can use the same standby group for command switch redundancy and HSRP redundancy. Use the **cluster standby-group** *HSRP-group-name* [**routing-redundancy**] global configuration command to enable the same HSRP standby group to be used for command switch and routing redundancy. If you create a cluster with the same HSRP standby group name without entering the **routing-redundancy** keyword, HSRP standby routing is disabled for the group.

This example shows how to bind standby group my_hsrp to the cluster and enable the same HSRP group to be used for command switch redundancy and router redundancy. The command can only be executed on the command switch. If the standby group name or number does not exist, or if the switch is a member switch, an error message appears.

```
Switch# configure terminal
Switch(config)# cluster standby-group my_hsrp routing-redundancy
Switch(config)# end
```

Displaying HSRP Configurations

From privileged EXEC mode, use this command to display HSRP settings:

show standby [*interface-id* [*group*]] [**brief**] [**detail**]

You can display HSRP information for the whole switch, for a specific interface, for an HSRP group, or for an HSRP group on an interface. You can also specify whether to display a concise overview of HSRP information or detailed HSRP information. The default display is **detail**. If there are a large number of HSRP groups, using the **show standby** command without qualifiers can result in an unwieldy display.

This is a an example of output from the **show standby** privileged EXEC command, displaying HSRP information for two standby groups (group 1 and group 100):

```
Switch# show standby
VLAN1 - Group 1
   Local state is Standby, priority 105, may preempt
   Hellotime 3 holdtime 10
   Next hello sent in 00:00:02.182
   Hot standby IP address is 10.0.0.1 configured
   Active router is 172.20.138.35 expires in 00:00:09
   Standby router is local
   Standby virtual mac address is 0000.0c07.ac01
   Name is bbb
VLAN1 - Group 100
   Local state is Active, priority 105, may preempt
   Hellotime 3 holdtime 10
   Next hello sent in 00:00:02.262
   Hot standby IP address is 172.20.138.51 configured
   Active router is local
   Standby router is unknown expired
   Standby virtual mac address is 0000.0c07.ac64
   Name is test
```

CHAPTER 24

Configuring IP Multicast Routing

IP multicasting is a more efficient way to use network resources, especially for bandwidth-intensive services such as audio and video. IP multicast allows a host (source) to send packets to a group of hosts (receivers) anywhere within the IP network by using a special form of IP address called the IP *multicast group address.* The sending host inserts the multicast group address into the IP destination address field of the packet, and IP multicast routers and multilayer switches forward incoming IP multicast packets out all interfaces that lead to members of the multicast group.

IP multicast addresses are assigned to the old class D address space by the Internet Assigned Number Authority (IANA). The high-order bits of a Class D address are 1110. Therefore, host group addresses can be in the range 224.0.0.0 through 239.255.255.255. The address 224.0.0.0 is guaranteed not to be assigned to any group. The address 224.0.0.1 is assigned to the all-hosts multicast group on a subnet. The address 224.0.0.2 is assigned to the all-multicast-routers group on a subnet.

Any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group receive the message. Membership in a multicast group is dynamic; hosts can join and leave at any time. There is no restriction on the location or number of members in a multicast group. A host can be a member of more than one multicast group at a time. How active a multicast group is and what members it has can vary from group to group and from time to time. A multicast group can be active for a long time, or it can be very short-lived. Membership in a group can constantly change. A group that has members can have no activity.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter describes how to configure IP multicast routing on your multilayer switch. To use this feature, you must have the enhanced multilayer software image (EMI) installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

This chapter consists of these sections:

- Cisco Implementation of IP Multicast Routing, page 24-2
- Configuring IP Multicast Routing, page 24-13
- Configuring Advanced PIM Features, page 24-28
- Configuring Optional IGMP Features, page 24-31
- Configuring Optional Multicast Routing Features, page 24-37

- Configuring Basic DVMRP Interoperability Features, page 24-43
- Configuring Advanced DVMRP Interoperability Features, page 24-50
- Monitoring and Maintaining IP Multicast Routing, page 24-57

For information on configuring the Multicast Source Discovery Protocol (MSDP), see Chapter 25, "Configuring MSDP."

Note When you are configuring multicast routing parameters for the switch, to allocate system resources to maximize the number of possible multicast routes allowed, you can use the **sdm prefer routing** global configuration command to set the Switch Database Management feature to the routing template. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-57.

Cisco Implementation of IP Multicast Routing

The Cisco IOS software supports these protocols to implement IP multicast routing:

- Internet Group Management Protocol (IGMP) is used among hosts on a LAN and the routers (and multilayer switches) on that LAN to track the multicast groups of which hosts are members.
- Protocol-Independent Multicast (PIM) protocol is used among routers and multilayer switches to track which multicast packets to forward to each other and to their directly connected LANs.
- Distance Vector Multicast Routing Protocol (DVMRP) is used on the multicast backbone of the Internet (MBONE). The Cisco IOS software supports PIM-to-DVMRP interaction.
- Cisco Group Management Protocol (CGMP) is used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP.

Figure 24-1 shows where these protocols operate within the IP multicast environment.

Figure 24-1 IP Multicast Routing Protocols

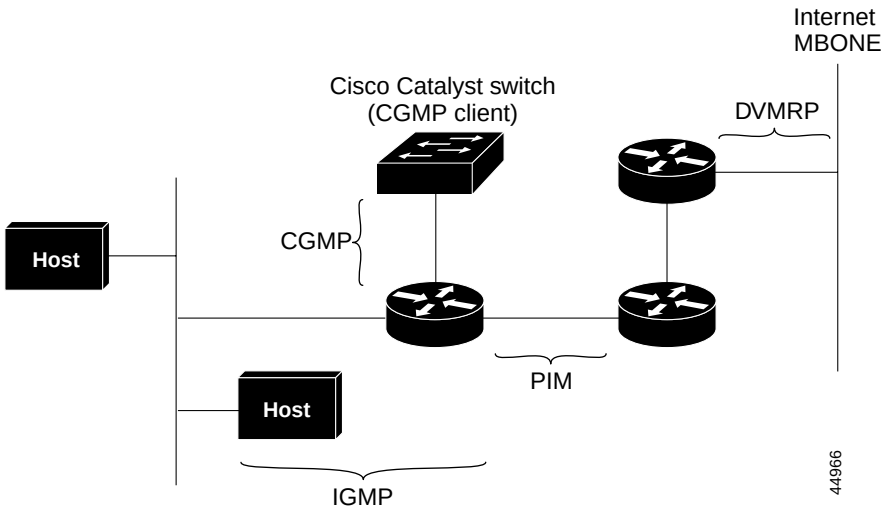

Understanding IGMP

To participate in IP multicasting, multicast hosts, routers, and multilayer switches must have IGMP operating. This protocol is the group membership protocol used by hosts to inform routers and multilayer switches of the existence of members on their directly connected networks and to allow them to send and receive multicast datagrams.

Multicast routers and switches learn about group membership when a host joining a new group sends an IGMP message to the group address declaring its membership.

Using the information obtained through IGMP, routers and switches maintain a list of multicast group memberships on a per-interface basis. A multicast group membership is active on an interface if at least one host on that interface has sent an IGMP join message to receive the multicast group traffic.

IGMP Version 1

Most IP stacks in hosts today still use IGMPv1. This version primarily uses a query-response model that allows the multicast router and multilayer switch to determine which multicast groups are active (have one or more hosts interested in a multicast group) on the local subnet. In this model, the router or switch acting as the IGMP querier periodically (every 60 seconds) multicasts an IGMPv1 membership query to the all-hosts multicast group (224.0.0.1) on the local subnet. All hosts enabled for multicasting listen for this address and receive the query. A host responds with an IGMPv1 membership report to receive multicast traffic for a specific group, and routers or switches on the subnet learn where active receivers are for the multicast groups.

A host can also join a multicast group by sending one or more unsolicited membership reports as shown in Figure 24-2. In this example, Host 3 sends an unsolicited report to receive traffic for multicast group 224.3.3.3 instead of waiting for the next membership query from Router 1.

A host leaves a multicast group by ceasing to process traffic for the multicast group and to respond to IGMP queries.

Figure 24-2 IGMPv1 Join Process

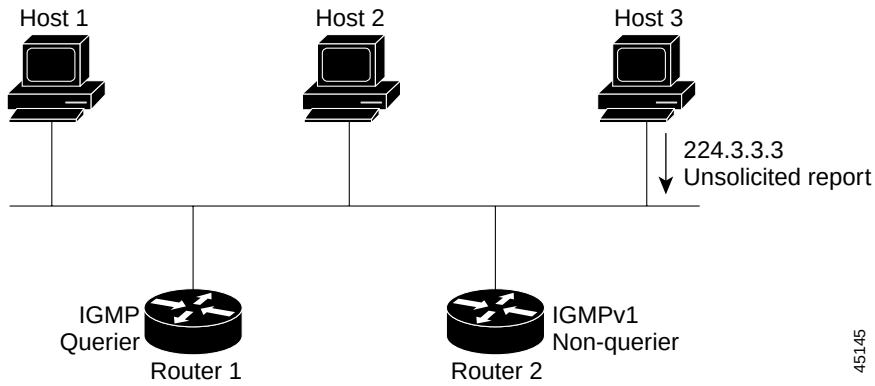

IGMPv1 relies on the Layer 3 IP multicast routing protocols (PIM, DVMRP, and so forth) to resolve which one of multiple multicast routers or multilayer switches on a subnet should be the querier. The query router sends IGMPv1 queries to determine which multicast groups are active (have one or more hosts sending unsolicited reports) on the local subnet. In general, a designated router is selected as the querier.

IGMP Version 2

IGMPv2 provides enhancements over IGMPv1. The query and membership report messages are identical to IGMPv1 message with two exceptions. The first difference is that the IGMPv2 query message is broken into two categories: general queries, which perform the same function as the IGMPv1 queries, and group-specific queries, which are queries directed to a single group. The second difference is that different type codes are used with IGMPv1 and IGMPv2 membership reports. IGMPv2 also includes new features:

- Querier election process—IGMPv2 routers or multilayer switches can elect the query router without having to rely on the multicast routing protocol to perform this process.

  As each IGMPv2 router or multilayer switch starts, it sends an IGMPv2 general query message to the all-host multicast group (224.0.0.1) with its interface address in the source IP address field of the message. Each IGMPv2 device compares the source IP address in the message with its own interface address, and the device with the lowest IP address on the subnet is elected as the querier.

- Maximum response time field—this field in the query message permits the query router to specify the maximum query-response time and controls the burstiness of the response process.

  This feature can be important when large numbers of groups are active on a subnet and you want to spread the responses over a longer period of time. However, increasing the maximum response timer value also increases the leave latency; the query router must now wait longer to make sure there are no more hosts for the group on the subnet.

- Group-specific query message—permits the query router to perform the query operation on a specific group instead of all groups.

- Leave group messages—provides hosts with a method of notifying routers and multilayer switches on the network that they are leaving a group as shown in Figure 24-3.

Figure 24-3 IGMPv2 Leave Process

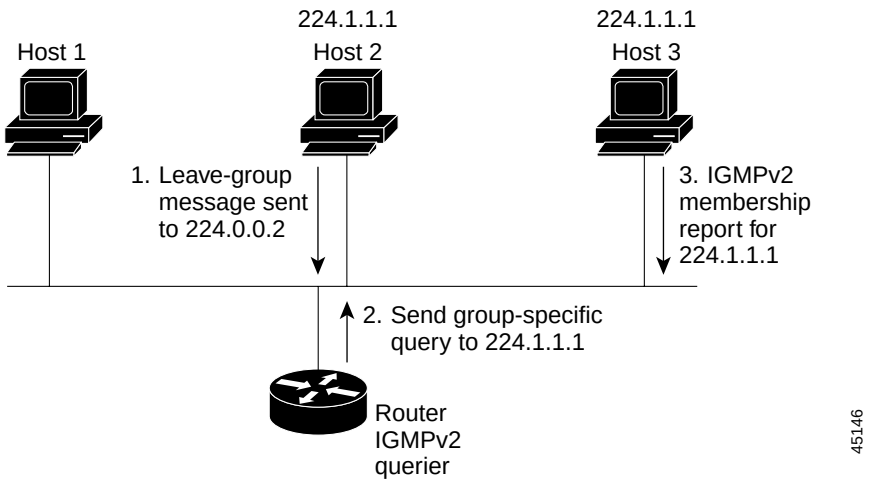

In this example, Hosts 2 and 3 are members of multicast group 224.1.1.1. Host 2 sends an IGMPv2 leave message to the all-multicast-routers group (224.0.0.2) to inform all routers and multilayer switches on the subnet that it is leaving the group. Router 1, the query router, receives the message, but because it keeps a list only of the group memberships that are active on a subnet and not individual hosts that are members, it sends a group-specific query to the target group (224.1.1.1) to determine whether any hosts remain for the group. Host 3 is still a member of multicast group 224.1.1.1 and receives the

group-specific query. It responds with an IGMPv2 membership report to inform Router 1 that a member is still present. When Router 1 receives the report, it keeps the group active on the subnet. If no response is received, the query router stops forwarding its traffic to the subnet.

Understanding PIM

PIM is called *protocol-independent*: regardless of the unicast routing protocols used to populate the unicast routing table, PIM uses this information to perform multicast forwarding instead of maintaining a separate multicast routing table.

PIM Versions

Two versions of PIM are supported in the IOS software. With PIM Version 1 (PIMv1), Cisco introduced support in IOS Release 11.1(6) for a new feature called Auto-RP. This proprietary feature eliminates the need to manually configure the rendezvous point (RP) information in every router and multilayer switch in the network. For more information, see the "Auto-RP" section on page 24-8.

Beginning with IOS Release 11.3, Cisco introduced support for PIM Version 2 (PIMv2) and its associated bootstrap router (BSR) capability. Like Auto-RP, the PIMv2 BSR mechanism eliminates the need to manually configure RP information in every router and multilayer switch in the network. For more information, see the "Bootstrap Router" section on page 24-8.

All systems using Cisco IOS Release 11.3(2)T or later start in PIMv2 mode by default. PIMv2 includes these improvements over PIMv1:

- A single, active RP exists per multicast group, with multiple backup RPs. This single RP compares to multiple active RPs for the same group in PIMv1.
- A BSR provides a fault-tolerant, automated RP discovery and distribution mechanism that enables routers and multilayer switches to dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. We strongly recommend sparse-dense mode, as opposed to either sparse mode or dense mode only.
- PIM join and prune messages have more flexible encoding for multiple address families.
- A more flexible hello packet format replaces the query packet to encode current and future capability options.
- Register messages to an RP specify whether they are sent by a border router or a designated router.
- PIM packets are no longer inside IGMP packets; they are standalone packets.

PIM Modes

PIM can operate in dense mode (DM), sparse mode (SM), or in sparse-dense mode (PIM DM-SM), which handles both sparse groups and dense groups at the same time.

PIM DM

In dense mode, a PIM DM router or multilayer switch assumes that all other routers or multilayer switches forward multicast packets for a group. If a PIM DM device receives a multicast packet and has no directly connected members or PIM neighbors present, a prune message is sent back to the source. Subsequent multicast packets are not flooded to this router or switch on this pruned branch. PIM DM builds source-based multicast distribution trees.

The simplest form of a multicast distribution tree is a source tree whose root is the source of the multicast traffic and whose branches form a spanning tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a *shortest-path tree* (SPT). A separate SPT exists for every individual source sending to each group. The special notation of (S,G) (pronounced S comma G) identifies an SPT where S is the IP address of the source and G is the multicast group address.

Figure 24-4 shows an example of SPT for group 224.1.1.1 rooted at the source, Host A, and connecting two receivers, Hosts B and C. The SPT notation for this group would be (194.1.1.1, 224.1.1.1).

Figure 24-4 Host A Shortest-Path Tree

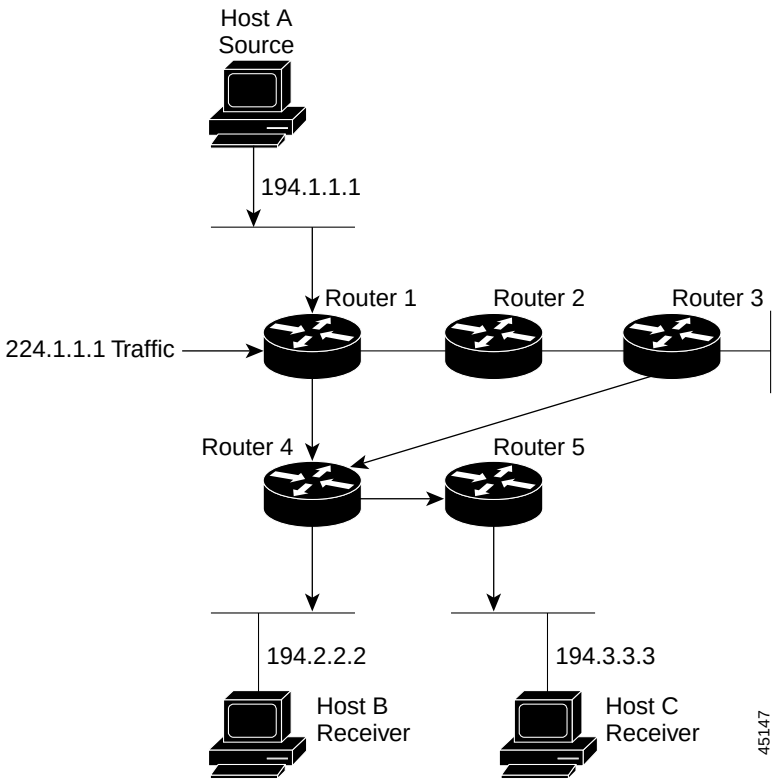

If Host B is also sending traffic to group 224.1.1.1 and Hosts A and C are receivers, then a separate (S,G) SPT would exist with the notation of (194.2.2.2, 224.1.1.1).

PIM DM employs only SPTs to deliver (S,G) multicast traffic by using a flood and prune method. It assumes that every subnet in the network has at least one receiver of the (S,G) multicast traffic, and therefore the traffic is flooded to all points in the network.

To avoid unnecessary consumption of network resources, PIM DM devices send prune messages up the source distribution tree to stop unwanted multicast traffic. Branches without receivers are pruned from the distribution tree, leaving only branches that contain receivers. Prunes have a timeout value associated with them, after which the PIM DM device puts the interface into the forwarding state and floods multicast traffic out the interface. When a new receiver on a previously pruned branch of the tree joins a multicast group, the PIM DM device detects the new receiver and immediately sends a graft message up the distribution tree toward the source. When the upstream PIM DM device receives the graft message, it immediately puts the interface on which the graft was received into the forwarding state so that the multicast traffic begins flowing to the receiver.

PIM SM

PIM SM uses shared trees and SPTs to distribute multicast traffic to multicast receivers in the network. In PIM SM, a router or multilayer switch assumes that other routers or switches do not forward multicast packets for a group, unless there is an explicit request for the traffic (join message). When a host joins a multicast group using IGMP, its directly connected PIM SM device sends PIM join messages toward the root, also known as the RP. This join message travels router-by-router toward the root, constructing a branch of the shared tree as it goes. The RP keeps track of multicast receivers; it also registers sources through register messages received from the source's first-hop router (*designated router* [DR]) to complete the shared tree path from the source to the receiver. The branches of the shared tree are maintained by periodic join refresh messages that the PIM SM devices send along the branch.

When using a shared tree, sources must send their traffic to the RP so that the traffic reaches all receivers. The special notation *,G, (pronounced star comma G) is used to represent the tree, where * means all sources and G represents the multicast group. Figure 24-5 shows a shared tree for group 224.2.2.2 with the RP located at Router 3. Multicast group traffic from source Hosts A and D travels to the RP (Router 3) and then down the shared tree to two receivers, Hosts B and C. Because all sources in the multicast group use a common shared tree, the special notation (*, 224.2.2.2) describes this shared tree.

Figure 24-5 Shared Distribution Tree

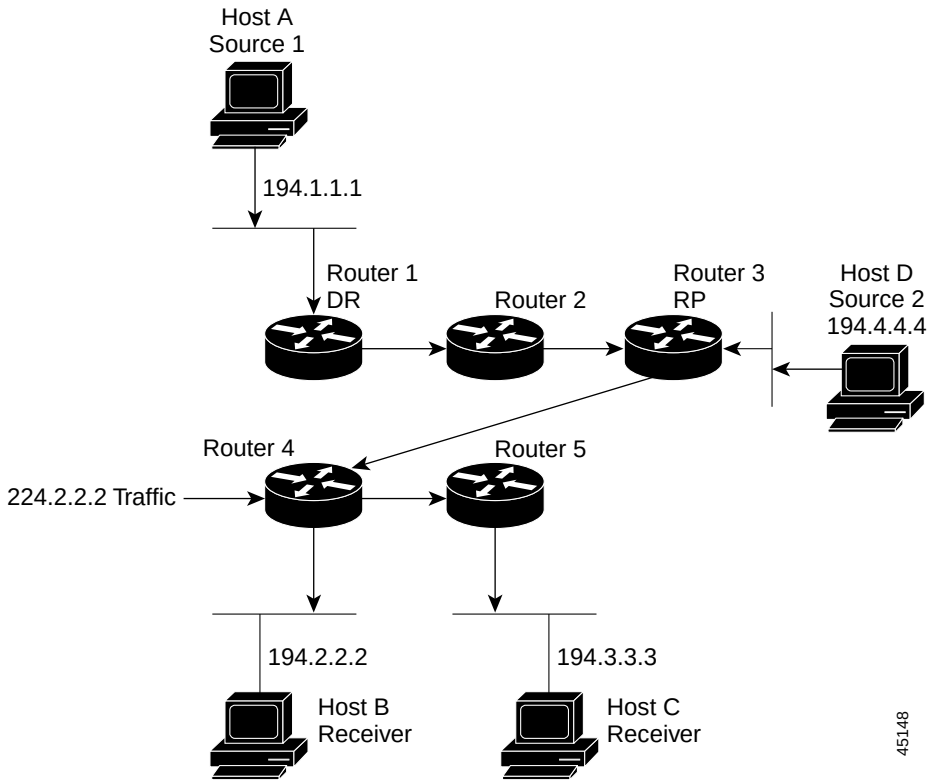

Note In addition to using the shared distribution tree, PIM SM can also use SPTs. By joining an SPT, multicast traffic is routed directly to the receivers without having to go through the RP, thereby reducing network latency and possible congestion at the RP. The disadvantage is that PIM SM devices must create and maintain (S,G) state entries in their routing tables along with the (S,G) SPT. This action consumes router resources.

Prune messages are sent up the distribution tree to prune multicast group traffic. This action permits branches of the shared tree or SPT that were created with explicit join messages to be torn down when they are no longer needed. For example, if a leaf router (a router without any downstream connections) detects that it no longer has any directly connected hosts (or downstream multicast routers) for a particular multicast group, it sends a prune message up the distribution tree to stop the flow of unwanted multicast traffic.

Auto-RP

This proprietary feature eliminates the need to manually configure the rendezvous point (RP) information in every router and multilayer switch in the network. For Auto-RP to work, you configure a Cisco router or multilayer switch as the mapping agent. It uses IP multicast to learn which routers or switches in the network are possible candidate RPs by joining the well-known Cisco-RP-announce multicast group (224.0.1.39) to receive candidate RP announcements. Candidate RPs send multicast RP-announce messages to a particular group or group range every 60 seconds (default) to announce their availability. Each RP-announce message contains a holdtime that tells the mapping agent how long the candidate RP announcement is valid. The default is 180 seconds.

Mapping agents listen to these candidate RP announcements and use the information to create entries in their Group-to-RP mapping caches. Only one mapping cache entry is created for any Group-to-RP range received, even if multiple candidate RPs are sending RP announcements for the same range. As the RP-announce messages arrive, the mapping agent selects the router or switch with the highest IP address as the active RP and stores this RP address in the Group-to-RP mapping cache.

Mapping agents multicast the contents of their Group-to-RP mapping cache in RP-discovery messages every 60 seconds (default) to the Cisco-RP-discovery multicast group (224.0.1.40), which all Cisco PIM routers and multilayer switches join to receive Group-to-RP mapping information. Thus, all routers and switches automatically discover which RP to use for the groups they support. The discovery messages also contain a holdtime, which defines how long the Group-to-RP mapping is valid. If a router or switch fails to receive RP-discovery messages and the Group-to-RP mapping information expires, it switches to a statically configured RP that was defined with the **ip pim rp-address** global configuration command. If no statically configured RP exists, the router or switch changes the group to dense-mode operation.

Multiple RPs serve different group ranges or serve as hot backups of each other.

Bootstrap Router

PIMv2 BSR is another method to distribute group-to-RP mapping information to all PIM routers and multilayer switches in the network. It eliminates the need to manually configure RP information in every router and switch in the network. However, instead of using IP multicast to distribute group-to-RP mapping information, BSR uses hop-by-hop flooding of special BSR messages to distribute the mapping information.

The BSR is elected from a set of candidate routers and switches in the domain that have been configured to function as BSRs. The election mechanism is similar to the root-bridge election mechanism used in bridged LANs. The BSR election is based on the BSR priority of the device contained in the BSR messages that are sent hop-by-hop through the network. Each BSR device examines the message and forwards out all interfaces only the message that has either a higher BSR priority than its BSR priority or the same BSR priority, but with a higher BSR IP address. Using this method, the BSR is elected.

The elected BSR sends BSR messages to the all-PIM-routers multicast group (224.0.0.13) with a TTL of 1. Neighboring PIMv2 routers or multilayer switches receive the BSR message and multicast it out all other interfaces (except the one on which it was received) with a TTL of 1. In this way, BSR messages

travel hop-by-hop throughout the PIM domain. Because BSR messages contain the IP address of the current BSR, the flooding mechanism allows candidate RPs to automatically learn which device is the elected BSR.

Candidate RPs send candidate RP advertisements showing the group range for which they are responsible directly to the BSR, which stores this information in its local candidate-RP cache. The BSR periodically advertises the contents of this cache in BSR messages to all other PIM devices in the domain. These messages travel hop-by-hop through the network to all routers and switches, which store the RP information in the BSR message in their local RP cache. The routers and switches select the same RP for a given group because they all use a common RP hashing algorithm.

Multicast Forwarding and Reverse Path Check

With unicast routing, routers and multilayer switches forward traffic through the network along a single path from the source to the destination host whose IP address appears in the destination address field of the IP packet. Each router and switch along the way makes a unicast forwarding decision, using the destination IP address in the packet, by looking up the destination address in the unicast routing table and forwarding the packet through the specified interface to the next hop toward the destination.

With multicasting, the source is sending traffic to an arbitrary group of hosts represented by a multicast group address in the destination address field of the IP packet. To determine whether to forward or drop an incoming multicast packet, the router or multilayer switch uses a reverse path forwarding (RPF) check on the packet as follows and shown in Figure 24-6:

1. The router or multilayer switch examines the source address of the arriving multicast packet to determine whether the packet arrived on an interface that is on the reverse path back to the source.
2. If the packet arrives on the interface leading back to the source, the RPF check is successful and the packet is forwarded to all interfaces in the outgoing interface list (which might not be all interfaces on the router).
3. If the RPF check fails, the packet is discarded.

Some multicast routing protocols, such as DVMRP, maintain a separate multicast routing table and use it for the RPF check. However, PIM uses the unicast routing table to perform the RPF check.

Figure 24-6 shows Gigabit Ethernet interface 0/2 receiving a multicast packet from source 151.10.3.21. A check of the routing table shows that the interface on the reverse path to the source is Gigabit Ethernet interface 0/1, not interface 0/2. Because the RPF check fails, the multilayer switch discards the packet. Another multicast packet from source 151.10.3.21 is received on interface 0/1, and the routing table shows this interface is on the reverse path to the source. Because the RPF check passes, the switch forwards the packet to all interfaces in the outgoing interface list.

Figure 24-6 RPF Check

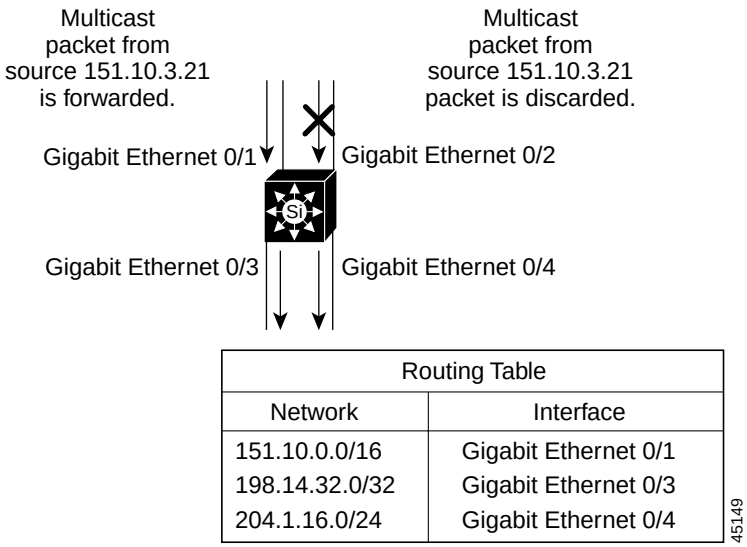

PIM uses both source trees and RP-rooted shared trees to forward datagrams (described in the "PIM DM" section on page 24-5 and the "PIM SM" section on page 24-7); the RPF check is performed differently for each:

- If a PIM router or multilayer switch has a source-tree state (that is, an (S,G) entry is present in the multicast routing table), it performs the RPF check against the IP address of the source of the multicast packet.
- If a PIM router or multilayer switch has a shared-tree state (and no explicit source-tree state), it performs the RPF check on the rendezvous point (RP) address (which is known when members join the group).

Sparse-mode PIM uses the RPF lookup function to determine where it needs to send joins and prunes:

- (S,G) joins (which are source-tree states) are sent toward the source.
- (*,G) joins (which are shared-tree states) are sent toward the RP.

DVMRP and dense-mode PIM use only source trees and use RPF as previously described.

Neighbor Discovery

PIM uses a neighbor discovery mechanism to establish PIM neighbor adjacencies. To establish adjacencies, a PIM router or multilayer switch sends PIM hello messages to the all-PIM-routers multicast group (224.0.0.13) on each of its multicast-enabled interfaces. The hello message contains a holdtime, which tells the receiver when the neighbor adjacency associated with the sender expires if no more PIM hello messages are received. (Keeping track of adjacencies is important for PIM DM operation for building the source distribution tree.)

PIM hello messages are also used to elect the DR for multi-access networks (Ethernet). The router or multilayer switch on the network with the highest IP address is the DR. With PIM DM operation, the DR has meaning only if IGMPv1 is in use; IGMPv1 does not have an IGMP querier election process, so the elected DR functions as the IGMP querier. In PIM SM operation, the DR is the router or switch that is directly connected to the multicast source. It sends PIM register messages to notify the RP that multicast traffic from a source needs to be forwarded down the shared tree.

Understanding DVMRP

Distance Vector Multicast Routing Protocol (DVMRP) is implemented in the equipment of many vendors and is based on the public-domain mrouted program. This protocol has been deployed in the multicast backbone (MBONE) and in other intradomain multicast networks.

Cisco routers and multilayer switches run PIM and can forward multicast packets to and receive from a DVMRP neighbor. It is also possible to propagate DVMRP routes into and through a PIM cloud. The Cisco IOS software propagates DVMRP routes and builds a separate database for these routes on each router and multilayer switch, but PIM uses this routing information to make the packet-forwarding decision. The Cisco IOS software does not implement the complete DVMRP. The Cisco IOS software supports dynamic discovery of DVMRP routers and can interoperate with them over traditional media (such as Ethernet and FDDI) or over DVMRP-specific tunnels.

DVMRP Neighbor Discovery

A DVMRP router learns about other DVMRP routers by periodically sending DVMRP probe messages to the all-DVMRP-routers multicast group (224.0.0.4). A second DVMRP router receiving the message adds the IP address of the first router that sent the probe to its internal list of DVMRP neighbors on the received interface and then sends its own probe message. This probe message contains all the addresses of neighboring DVMRP routers in its neighbor list, including the address of the first router. When the first DVMRP router receives a probe with its own address listed in the neighbor list, a two-way adjacency is formed between itself and the neighbor that sent the probe.

DVMRP Route Table

DVMRP neighbors build a route table by periodically exchanging source network routing information in route-report messages. These messages contain entries that advertise a source network with a mask and a hop count that is used as the routing metric. The routing information stored in the DVMRP routing table is separate from the unicast routing table and is used to build a source distribution tree and to perform multicast forward using reverse-path forwarding (RPF).

DVMRP Source Distribution Tree

DVMRP is a dense-mode protocol and builds a parent-child database using a constrained multicast model to build a forwarding tree rooted at the source of the multicast packets. Multicast packets are initially flooded down this source tree. If redundant paths are on the source tree, packets are not forwarded along those paths. Forwarding occurs until prune messages are received on those parent-child links, which further constrain the broadcast of multicast packets. DVMRP supports a reliable graft and graft-ack mechanism that grafts previously pruned branches of a tree. The graft-ack messages are sent by the upstream router in response to received graft messages, preventing the loss of a graft message because of congestion.

Understanding CGMP

This software release provides CGMP-server support on your multilayer switches; no client-side functionality is provided. The multilayer switch serves as a CGMP server for devices that do not support IGMP snooping but have CGMP-client functionality.

CGMP is a protocol used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP. CGMP permits Layer 2 group membership information to be communicated from the CGMP server to the switch, which can learn on which ports multicast members reside instead of flooding multicast traffic to all switch ports. (IGMP snooping is another method to constrain the flooding of multicast packets. For more information, see Chapter 11, "Configuring IGMP Snooping and MVR.")

CGMP is necessary because the Layer 2 switch cannot distinguish between IP multicast data packets and IGMP report messages, which are both at the MAC-level and are addressed to the same group address.

Joining a Group with CGMP

Hosts connected to a Layer 2 Catalyst switch can join a multicast group by sending an unsolicited IGMP membership report message to the target group (224.1.2.3) as shown in Figure 24-7. Because LAN switches operate at Layer 2 and understand only MAC addresses, the source and destination fields of the frame contain 48-bit MAC addresses for Host 3 (0080.c7a2.1093) and MAC-address equivalent of the multicast group address (0100.5e01.0203).

The IGMP membership report is received by the Layer 2 switch and forwarded to the CGMP server for normal IGMP processing. The CGMP server, which must have CGMP enabled on the interface connected to the Layer 2 switch, receives the membership report and translates the report into a CGMP join message. It sends the CGMP join message to the switch through the well-known CGMP multicast MAC address (0x0100.0cdd.dddd). When the Layer 2 switch receives the join message, it updates its forwarding table to include the MAC-equivalent of the group destination address and the applicable input and output switch ports.

Figure 24-7 Host Joining a Group Using CGMP

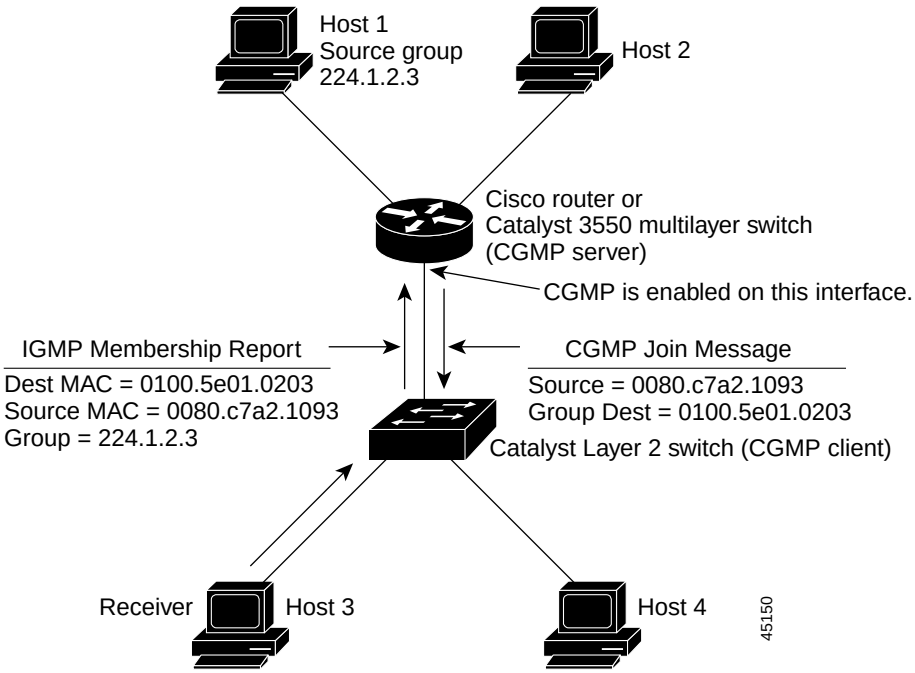

Leaving a Group with CGMP

When an IGMPv2 host leaves a group, it can send an IGMP leave group message to the all-multicast-routers group (224.0.0.2). The CGMP server translates this leave group message into a CGMP leave message and sends it to the switch.

To expedite a host on a LAN leaving a multicast group, some Layer 2 Catalyst switch software offers the CGMP Fast-Leave feature, which allows the switch to perform IGMPv2 leave processing locally without involving the CGMP server and accelerates the removal of unused CGMP groups. The host sends the leave group message to the all-multicast-routers group (224.0.0.2). The Layer 2 switch processes it and does not forward it to the CGMP server. The Layer 2 switch sends an IGMP general query message on the port where the leave message was received to determine if there are remaining members for the group on the port. If no response is received, the Layer 2 switch sends an IGMP leave message to the CGMP server, which sends a group-specific query to the multicast group to see if there are any remaining members in the group. If there is no response, the CGMP server updates its multicast routing table and sends a CGMP delete group message to the Layer 2 switch, which updates its routing table.

Configuring IP Multicast Routing

This section describes how to configure IP multicast routing. It contains this configuration information:

- Default Multicast Routing Configuration, page 24-13
- Multicast Routing Configuration Guidelines, page 24-14
- Configuring Basic Multicast Routing, page 24-15 (required procedure)
- Configuring a Rendezvous Point, page 24-17 (required for spare-mode or sparse-dense-mode operation)
- Using Auto-RP and a BSR, page 24-27
- Monitoring the RP Mapping Information, page 24-27
- Troubleshooting PIMv1 and PIMv2 Interoperability Problems, page 24-28

Default Multicast Routing Configuration

Table 24-1 shows the default multicast routing configuration.

Table 24-1 Default Multicast Routing Configuration

Feature	Default Setting
Multicast routing	Disabled on all interfaces.
PIM version	Version 2 (for devices running IOS Release 11.3(2)T or later).
PIM mode	No mode is defined.
PIM RP address	None configured.
PIM domain border	Disabled.
PIM multicast boundary	None.
Candidate BSRs	Disabled.

Table 24-1 Default Multicast Routing Configuration (continued)

Feature	Default Setting
Candidate RPs	Disabled.
Shortest-path tree threshold rate	0 kbps.
PIM router query message interval	30 seconds.

Multicast Routing Configuration Guidelines

To avoid misconfiguring multicast routing on your multilayer switch, review the information in these sections:

- PIMv1 and PIMv2 Interoperability, page 24-14
- Auto-RP and BSR Configuration Guidelines, page 24-15

PIMv1 and PIMv2 Interoperability

The Cisco PIMv2 implementation allows interoperability and transition between Version 1 and Version 2, although there might be some minor problems.

You can upgrade to PIMv2 incrementally. PIM Versions 1 and 2 can be configured on different routers and multilayer switches within one network. Internally, all routers and multilayer switches on a shared media network must run the same PIM version. Therefore, if a PIMv2 device detects a PIMv1 device, the Version 2 device downgrades itself to Version 1 until all Version 1 devices have been shut down or upgraded.

PIMv2 uses the BSR to discover and announce RP-set information for each group prefix to all the routers and multilayer switches in a PIM domain. PIMv1, together with the Auto-RP feature, can perform the same tasks as the PIMv2 BSR. However, Auto-RP is a standalone protocol, separate from PIMv1, and is a proprietary Cisco protocol. PIMv2 is a standards track protocol in the IETF. We recommend that you use PIMv2. The BSR mechanism interoperates with Auto-RP on Cisco routers and multilayer switches. For more information, see the "Auto-RP and BSR Configuration Guidelines" section on page 24-15.

When PIMv2 devices interoperate with PIMv1 devices, Auto-RP should have already been deployed. A PIMv2 BSR that is also an Auto-RP mapping agent automatically advertises the RP elected by Auto-RP. That is, Auto-RP sets its single RP on every router or multilayer switch in the group. Not all routers and switches in the domain use the PIMv2 hash function to select multiple RPs.

Dense-mode groups in a mixed PIMv1 and PIMv2 region need no special configuration; they automatically interoperate.

Sparse-mode groups in a mixed PIMv1 and PIMv2 region are possible because the Auto-RP feature in PIMv1 interoperates with the PIMv2 RP feature. Although all PIMv2 devices can also use PIMv1, we recommend that the RPs be upgraded to PIMv2 (or at least upgraded to PIMv1 in the Cisco IOS Release 11.3 software). To ease the transition to PIMv2, we have these recommendations:

- Use Auto-RP throughout the region.
- Configure sparse-dense mode throughout the region.

If Auto-RP is not already configured in the PIMv1 regions, configure Auto-RP. For more information, see the "Configuring Auto-RP" section on page 24-18.

Auto-RP and BSR Configuration Guidelines

There are two approaches to using PIMv2. You can use Version 2 exclusively in your network or migrate to Version 2 by employing a mixed PIM version environment.

- If your network is all Cisco routers and multilayer switches, you can use either Auto-RP or BSR.
- If you have non-Cisco routers in your network, you must use BSR.
- If you have Cisco PIMv1 and PIMv2 routers and multilayer switches and non-Cisco routers, you must use both Auto-RP and BSR.
- Because bootstrap messages are sent hop-by-hop, a PIMv1 device prevents these messages from reaching all routers and multilayer switches in your network. Therefore, if your network has a PIMv1 device in it and only Cisco routers and multilayer switches, it is best to use Auto-RP.
- If you have a network that includes non-Cisco routers, configure the Auto-RP mapping agent and the BSR on a Cisco PIMv2 router or multilayer switch. Ensure that no PIMv1 device is on the path between the BSR and a non-Cisco PIMv2 router.
- If you have non-Cisco PIMv2 routers that need to interoperate with Cisco PIMv1 routers and multilayer switches, both Auto-RP and a BSR are required. We recommend that a Cisco PIMv2 device be both the Auto-RP mapping agent and the BSR. For more information, see the "Using Auto-RP and a BSR" section on page 24-27.

Configuring Basic Multicast Routing

You must enable IP multicast routing and configure the PIM version and PIM mode so that the IOS software can forward multicast packets and determine how the multilayer switch populates its multicast routing table.

You can configure an interface to be in PIM dense mode, sparse mode, or sparse-dense mode. The mode determines how the switch populates its multicast routing table and how it forwards multicast packets it receives from its directly connected LANs. You must enable PIM in one of these modes for an interface to perform IP multicast routing. Enabling PIM on an interface also enables IGMP operation on that interface.

By default, multicast routing is disabled, and there is no default mode setting. The following procedure is required.

Beginning in privileged EXEC mode, follow these steps to enable IP multicasting and a PIM mode on your multilayer switch:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip multicast-routing**	Enable IP multicast forwarding.

	Command	Purpose
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the Layer 3 interface on which you want to enable multicast routing. The specified interface must be one of the following: • A routed port: a physical port that has been configured as a Layer 3 port by entering the **no switchport** interface configuration command. • An SVI: a VLAN interface created by using the **interface vlan** *vlan-id* global configuration command. These ports must have IP addresses assigned to them. For more information, see the "Configuring Layer 3 Interfaces" section on page 8-22.
Step 4	**ip pim version** [**1** | **2**]	Configure the PIM version on the interface. By default, Version 2 is enabled and is the recommended setting. **Note** All IP multicast-capable Cisco PIM routers using IOS Release 11.3(2)T or later start in PIMv2 by default. An interface in PIMv2 mode automatically downgrades to PIMv1 mode if that interface has a PIMv1 neighbor. The interface returns to Version 2 mode after all Version 1 neighbors are shut down or upgraded. For more information, see the "PIMv1 and PIMv2 Interoperability" section on page 24-14.
Step 5	**ip pim** {**dense-mode** | **sparse-mode** | **sparse-dense-mode**}	Enable a PIM mode on the interface. By default, no mode is configured. The keywords have these meanings: • **dense-mode**—Enables dense mode of operation. • **sparse-mode**—Enables sparse mode of operation. • **sparse-dense-mode**—Causes the interface to be treated in the mode in which the group belongs. Sparse-dense-mode is the recommended setting. **Note** If you are use sparse-mode or sparse-dense mode, you must also configure an RP. For more information, see the "Configuring a Rendezvous Point" section on page 24-17.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable multicasting, use the **no ip multicast-routing** global configuration command. To return to the default PIM version, use the **no ip pim version** interface configuration command. To disable PIM on an interface, use the **no ip pim** interface configuration command.

Configuring a Rendezvous Point

If you have configured PIM SM or PIM SM-DM, you must configure an RP for the multicast group. You can use several methods, as described in the next sections:

- Manually Assigning an RP to Multicast Groups, page 24-17
- Configuring Auto-RP, page 24-18 (a standalone, Cisco-proprietary protocol separate from PIMv1)
- Configuring PIMv2 BSR, page 24-22 (a standards track protocol in the Internet Engineering Task Force (IETF)

You can use Auto-RP, BSR, or a combination of both, depending on the PIM version you are running and the types of routers in your network. For more information, see the "PIMv1 and PIMv2 Interoperability" section on page 24-14 and the "Auto-RP and BSR Configuration Guidelines" section on page 24-15.

Manually Assigning an RP to Multicast Groups

This section explains how to manually configure an RP. If the RP for a group is learned through a dynamic mechanism (such as Auto-RP or BSR), you need not perform this task for that RP.

Senders of multicast traffic announce their existence through register messages received from the source's first-hop router (designated router) and forwarded to the RP. Receivers of multicast packets use RPs to join a multicast group by using explicit join messages. RPs are not members of the multicast group; rather, they serve as a *meeting place* for multicast sources and group members.

Beginning in privileged EXEC mode, follow these steps to manually configure the address of the RP:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip pim rp-address** *ip-address* [*access-list-number*] [**override**]	Configure the address of a PIM RP. By default, no PIM RP address is configured. You must configure the IP address of RPs on all routers and multilayer switches (including the RP). If there is no RP configured for a group, the multilayer switch treats the group as dense, using the dense-mode PIM techniques. A PIM device can use multiple RPs, but only one per group. • For *ip-address*, enter the unicast address of the RP in dotted-decimal notation. • (Optional) For *access-list-number*, enter an IP standard access list number from 1 to 99. If no access list is configured, the RP is used for all groups. • (Optional) The **override** keyword means that if there is a conflict between the RP configured with this command and one learned by Auto-RP or BSR, the RP configured with this command prevails.

	Command	Purpose
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the multicast group address for which the RP should be used. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove an RP address, use the **no ip pim rp-address** *ip-address* [*access-list-number*] [**override**] global configuration command.

This example shows how to configure the address of the RP to 147.106.6.22 for multicast group 225.2.2.2 only:

```
Switch(config)# access-list 1 permit 225.2.2.2 0.0.0.0
Switch(config)# ip pim rp-address 147.106.6.22 1
```

Configuring Auto-RP

Auto-RP uses IP multicast to automate the distribution of group-to-RP mappings to all Cisco routers and multilayer switches in a PIM network. It has these benefits:

- It is easy to use multiple RPs within a network to serve different group ranges.
- It allows load splitting among different RPs and arrangement of RPs according to the location of group participants.
- It avoids inconsistent, manual RP configurations on every router and multilayer switch in a PIM network, which can cause connectivity problems.

Note If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP, you must manually configure an RP as described in the "Manually Assigning an RP to Multicast Groups" section on page 24-17.

Note If routed interfaces are configured in sparse mode, Auto-RP can still be used if all devices are configured with a manual RP address for the Auto-RP groups.

This section contains this configuration information:

- Setting up Auto-RP in a New Internetwork, page 24-19
- Adding Auto-RP to an Existing Sparse-Mode Cloud, page 24-19
- Preventing Join Messages to False RPs, page 24-20
- Preventing Candidate RP Spoofing, page 24-21

For overview information, see the "Auto-RP" section on page 24-8.

Setting up Auto-RP in a New Internetwork

If you are setting up Auto-RP in a new internetwork, you do not need a default RP because you configure all the interfaces for sparse-dense mode. Follow the process described in the section "Adding Auto-RP to an Existing Sparse-Mode Cloud" section on page 24-19. However, skip Step 3 to configure a PIM router as the RP for the local group.

Adding Auto-RP to an Existing Sparse-Mode Cloud

This section contains some suggestions for the initial deployment of Auto-RP into an existing sparse-mode cloud to minimize disruption of the existing multicast infrastructure.

Beginning in privileged EXEC mode, follow these steps to deploy Auto-RP in an existing sparse-mode cloud:

	Command	Purpose
Step 1	**show running-config**	Verify that a default RP is already configured on all PIM devices and the RP in the sparse-mode network. This step is not required for spare-dense-mode environments. The selected RP should have good connectivity and be available across the network. Use this RP for the global groups (for example 224.x.x.x and other global groups). Do not reconfigure the group address range that this RP serves. RPs dynamically discovered through Auto-RP take precedence over statically configured RPs. Assume that it is desirable to use a second RP for the local groups.
Step 2	**configure terminal**	Enter global configuration mode.
Step 3	**ip pim send-rp-announce** *interface-id* **scope** *ttl* **group-list** *access-list-number* **interval** *seconds*	Configure another PIM device to be the candidate RP for local groups. • For *interface-id*, enter the interface type and number that identifies the RP address. Valid interfaces include physical ports, port channels, and VLANs. • For **scope** *ttl*, specify the time-to-live value in hops. Enter a hop count that is high enough so that the RP-announce messages reach all mapping agents in the network. There is no default setting. The range is 1 to 255. • For **group-list** *access-list-number*, enter an IP standard access list number from 1 to 99. If no access list is configured, the RP is used for all groups. • For **interval** *seconds*, specify how often the announcement messages must be sent. The default is 60 seconds. The range is 1 to 16383.

	Command	Purpose
Step 4	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 3. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the multicast group address range for which the RP should be used. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 5	**ip pim send-rp-discovery scope** *ttl*	Find a multilayer switch whose connectivity is not likely to be interrupted, and assign it the role of RP-mapping agent. For **scope** *ttl*, specify the time-to-live value in hops to limit the RP discovery packets. All devices within the hop count from the source device receive the Auto-RP discovery messages. These messages tell other devices which group-to-RP mapping to use to avoid conflicts (such as overlapping group-to-RP ranges). There is no default setting. The range is 1 to 255.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show running-config**	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the PIM device configured as the candidate RP, use the **no ip pim send-rp-announce** global configuration command. To remove the multilayer switch as the RP-mapping agent, use the **no ip pim send-rp-discovery** global configuration command.

This example shows how to send RP announcements out all PIM-enabled interfaces for a maximum of 31 hops. The IP address of Gigabit Ethernet interface 0/1 is the RP. Access list 5 describes the group for which this multilayer switch serves as RP:

```
Switch(config)# ip pim send-rp-announce gigabitethernet0/1 scope 31 group-list 5
Switch(config)# access-list 5 permit 224.0.0.0 15.255.255.255
```

Preventing Join Messages to False RPs

Determine whether the **ip pim accept-rp** command was previously configured throughout the network by using the **show running-config** privileged EXEC command. If the **ip pim accept-rp** command is not configured on any device, this problem can be addressed later. In those routers or multilayer switches already configured with the **ip pim accept-rp** command, you must enter the command again to accept the newly advertised RP.

To accept all RPs advertised with Auto-RP and reject all other RPs by default, use the **ip pim accept-rp auto-rp** global configuration command.

If all interfaces are in sparse mode, use a default-configured RP to support the two well-known groups 224.0.1.39 and 224.0.1.40. Auto-RP uses these two well-known groups to collect and distribute RP-mapping information. When this is the case and the **ip pim accept-rp auto-rp** command is configured, another **ip pim accept-rp** command accepting the RP must be configured as follows:

```
Switch(config)# ip pim accept-rp 172.10.20.1 1
Switch(config)# access-list 1 permit 224.0.1.39
Switch(config)# access-list 1 permit 224.0.1.40
```

Preventing Candidate RP Spoofing

You can add configuration commands to the mapping agents to prevent a maliciously configured router from masquerading as a candidate RP and causing problems.

Beginning in privileged EXEC mode, follow these steps to filter incoming RP announcement messages:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip pim rp-announce-filter rp-list** *access-list-number* **group-list** *access-list-number*	Filter incoming RP announcement messages. Enter this command on each mapping agent in the network. Without this command, all incoming RP-announce messages are accepted by default. For **rp-list** *access-list-number*, configure an access list of candidate RP addresses that, if permitted, is accepted for the group ranges supplied in the **group-list** *access-list-number* variable. If this variable is omitted, the filter applies to all multicast groups. If more than one mapping agent is used, the filters must be consistent across all mapping agents to ensure that no conflicts occur in the Group-to-RP mapping information.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • Create an access list that specifies from which routers and multilayer switches the mapping agent accepts candidate RP announcements (rp-list ACL). • Create an access list that specifies the range of multicast groups from which to accept or deny (group-list ACL). • For *source*, enter the multicast group address range for which the RP should be used. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a filter on incoming RP announcement messages, use the **no ip pim rp-announce-filter rp-list** *access-list-number* **group-list** *access-list-number* global configuration command.

This example shows a sample configuration on an Auto-RP mapping agent that is used to prevent candidate RP announcements from being accepted from unauthorized candidate RPs:

```
Switch(config)# ip pim rp-announce-filter rp-list 10 group-list 20
Switch(config)# access-list 10 permit host 172.16.5.1
Switch(config)# access-list 10 permit host 172.16.2.1
Switch(config)# access-list 20 deny 239.0.0.0 0.0.255.255
Switch(config)# access-list 20 permit 224.0.0.0 15.255.255.255
```

In this example, the mapping agent accepts candidate RP announcements from only two devices, 172.16.5.1 and 172.16.2.1. The mapping agent accepts candidate RP announcements from these two devices only for multicast groups that fall in the group range of 224.0.0.0 to 239.255.255.255. The mapping agent does not accept candidate RP announcements from any other devices in the network. Furthermore, the mapping agent does not accept candidate RP announcements from 172.16.5.1 or 172.16.2.1 if the announcements are for any groups in the 239.0.0.0 through 239.255.255.255 range. This range is the administratively scoped address range.

Configuring PIMv2 BSR

BSR automates the distribution of group-to-RP mappings to all routers and multilayer switches in a PIMv2 network. It eliminates the need to manually configure RP information in every device in the network. However, instead of using IP multicast to distribute group-to-RP mapping information, BSR uses hop-by-hop flooding of special BSR messages to distribute the mapping information. For overview information, see the "Bootstrap Router" section on page 24-8.

This section describes how to set up BSR in your PIMv2 network. It contains this configuration information:

- Defining the PIM Domain Border, page 24-22
- Defining the IP Multicast Boundary, page 24-24
- Configuring Candidate BSRs, page 24-25
- Configuring Candidate RPs, page 24-26

Defining the PIM Domain Border

As IP multicast becomes more widespread, the chances of one PIMv2 domain bordering another PIMv2 domain is increasing. Because these two domains probably do not share the same set of RPs, BSR, candidate RPs, and candidate BSRs, you need to constrain PIMv2 BSR messages from flowing into or out of the domain. Allowing these messages to leak across the domain borders could adversely affect the normal BSR election mechanism and elect a single BSR across all bordering domains and co-mingle candidate RP advertisements, resulting in the election of RPs in the wrong domain.

Beginning in privileged EXEC mode, follow these steps to define the PIM domain border:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip pim bsr-border**	Define a PIM bootstrap message boundary for the PIM domain. Enter this command on each interface that connects to other bordering PIM domains. This command instructs the multilayer switch to neither send or receive PIMv2 BSR messages on this interface as shown in Figure 24-8.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the PIM border, use the **no ip pim bsr-border** interface configuration command.

Figure 24-8 Constraining PIMv2 BSR Messages

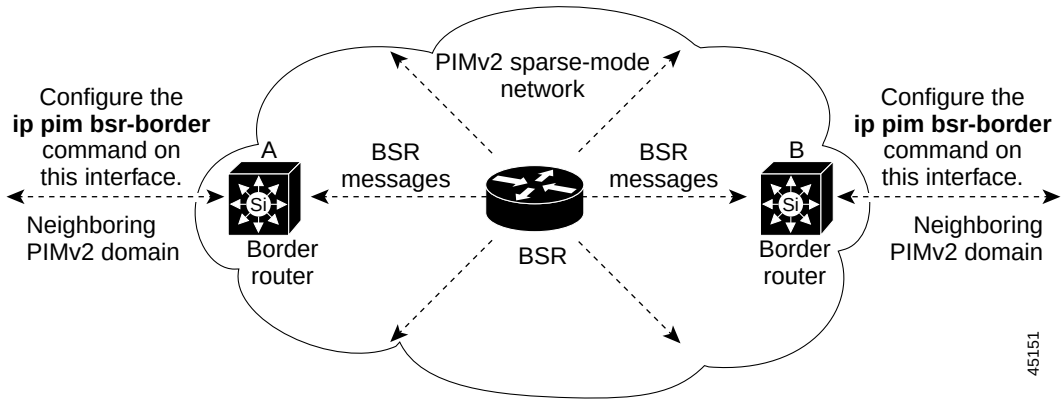

Defining the IP Multicast Boundary

You define a multicast boundary to prevent Auto-RP messages from entering the PIM domain. You create an access list to deny packets destined for 224.0.1.39 and 224.0.1.40, which carry Auto-RP information.

Beginning in privileged EXEC mode, follow these steps to define a multicast boundary:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* **deny** *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. • For *source*, enter multicast addresses 224.0.1.39 and 224.0.1.40, which carry Auto-RP information. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 4	**ip multicast boundary** *access-list-number*	Configure the boundary, specifying the access list you created in Step 2.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the boundary, use the **no ip multicast boundary** interface configuration command.

This example shows a portion of an IP multicast boundary configuration that denies Auto-RP information:

```
Switch(config)# access-list 1 deny 224.0.1.39
Switch(config)# access-list 1 deny 224.0.1.40
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip multicast boundary 1
```

Configuring Candidate BSRs

You can configure one or more candidate BSRs. The devices serving as candidate BSRs should have good connectivity to other devices and be in the backbone portion of the network.

Beginning in privileged EXEC mode, follow these steps to configure your multilayer switch as a candidate BSR:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip pim bsr-candidate** *interface-id hash-mask-length* [*priority*]	Configure your multilayer switch to be a candidate BSR. • For *interface-id*, enter the interface type and number on this switch from which the BSR address is derived to make it a candidate. This interface must be enabled with PIM. Valid interfaces include physical ports, port channels, and VLANs. • For *hash-mask-length*, specify the mask length (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash correspond to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. • (Optional) For *priority*, enter a number from 0 to 255. The BSR with the larger priority is preferred. If the priority values are the same, the device with the highest IP address is selected as the BSR. The default is 0.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove this device as a candidate BSR, use the **no ip pim bsr-candidate** global configuration command.

This example shows how to configure a candidate BSR, which uses the IP address 172.21.24.18 on Gigabit Ethernet interface 0/2 as the advertised BSR address, uses 30 bits as the hash-mask-length, and has a priority of 10.

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip address 172.21.24.18 255.255.255.0
Switch(config-if)# ip pim sparse-dense-mode
Switch(config-if)# ip pim bsr-candidate gigabitethernet0/2 30 10
```

Configuring Candidate RPs

You can configure one or more candidate RPs. Similar to BSRs, the RPs should also have good connectivity to other devices and be in the backbone portion of the network. An RP can serve the entire IP multicast address space or a portion of it. Candidate RPs send candidate RP advertisements to the BSR. When deciding which devices should be RPs, consider these options:

- In a network of Cisco routers and multilayer switches where only Auto-RP is used, any device can be configured as an RP.
- In a network that includes only Cisco PIMv2 routers and multilayer switches and with routers from other vendors, any device can be used as an RP.
- In a network of Cisco PIMv1 routers, Cisco PIMv2 routers, and routers from other vendors, configure only Cisco PIMv2 routers and multilayer switches as RPs.

Beginning in privileged EXEC mode, follow these steps to configure your multilayer switch to advertise itself as a PIMv2 candidate RP to the BSR:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip pim rp-candidate** *interface-id* [**group-list** *access-list-number*]	Configure your multilayer switch to be a candidate RP. • For *interface-id*, enter the interface type and number whose associated IP address is advertised as a candidate RP address. Valid interfaces include physical ports, port channels, and VLANs. • (Optional) For **group-list** *access-list-number*, enter an IP standard access list number from 1 to 99. If no group-list is specified, the multilayer switch is a candidate RP for all groups.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, enter the access list number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the number of the network or host from which the packet is being sent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove this device as a candidate RP, use the **no ip pim rp-candidate** global configuration command.

This example shows how to configure the multilayer switch to advertise itself as a candidate RP to the BSR in its PIM domain. Standard access list number 4 specifies the group prefix associated with the RP that has the address identified by Gigabit Ethernet interface 0/2. That RP is responsible for the groups with the prefix 239.

```
Switch(config)# ip pim rp-candidate gigabitethernet0/2 group-list 4
Switch(config)# access-list 4 permit 239.0.0.0 0.255.255.255
```

Using Auto-RP and a BSR

If you have non-Cisco PIMv2 routers that need to interoperate with Cisco PIMv1 routers and multilayer switches, both Auto-RP and a BSR are required. We recommend that a Cisco PIMv2 router or multilayer switch be both the Auto-RP mapping agent and the BSR.

If you must have one or more BSRs, we have these recommendations:

- Configure the candidate BSRs as the RP-mapping agents for Auto-RP. For more information, see the "Configuring Auto-RP" section on page 24-18 and the "Configuring Candidate BSRs" section on page 24-25.
- For group prefixes advertised through Auto-RP, the PIMv2 BSR mechanism should not advertise a subrange of these group prefixes served by a different set of RPs. In a mixed PIMv1 and PIMv2 domain, have backup RPs serve the same group prefixes. This prevents the PIMv2 DRs from selecting a different RP from those PIMv1 DRs, due to the longest match lookup in the RP-mapping database.

Beginning in privileged EXEC mode, follow these steps to verify the consistency of group-to-RP mappings:

	Command	Purpose
Step 1	**show ip pim rp** [[*group-name* | *group-address*] | **mapping**]	On any Cisco device, display the available RP mappings. • (Optional) For *group-name*, specify the name of the group about which to display RPs. • (Optional) For *group-address*, specify the address of the group about which to display RPs. • (Optional) Use the **mapping** keyword to display all group-to-RP mappings of which the Cisco device is aware (either configured or learned from Auto-RP).
Step 2	**show ip pim rp-hash** *group*	On a PIMv2 router or multilayer switch, confirm that the same RP is the one that a PIMv1 system chooses. For *group*, enter the group address for which to display RP information.

Monitoring the RP Mapping Information

To monitor the RP mapping information, use these commands in privileged EXEC mode:

- **show ip pim bsr** displays information about the currently elected BSR.
- **show ip pim rp-hash** *group* displays the RP that was selected for the specified group.
- **show ip pim rp** [*group-name* | *group-address* | **mapping**] displays how the multilayer switch learns of the RP (through the BSR or the Auto-RP mechanism).

Troubleshooting PIMv1 and PIMv2 Interoperability Problems

When debugging interoperability problems between PIMv1 and PIMv2, check these in the order shown:

1. Verify RP mapping with the **show ip pim rp-hash** privileged EXEC command, making sure that all systems agree on the same RP for the same group.
2. Verify interoperability between different versions of DRs and RPs. Make sure the RPs are interacting with the DRs properly (by responding with register-stops and forwarding decapsulated data packets from registers).

Configuring Advanced PIM Features

This section contains the optional advanced PIM features and this configuration information:

- Understanding PIM Shared Tree and Source Tree, page 24-28
- Delaying the Use of PIM Shortest-Path Tree, page 24-29
- Modifying the PIM Router-Query Message Interval, page 24-30

Understanding PIM Shared Tree and Source Tree

By default, members of a group receive data from senders to the group across a single data-distribution tree rooted at the RP. Figure 24-9 shows this type of shared-distribution tree. Data from senders is delivered to the RP for distribution to group members joined to the shared tree.

Figure 24-9 Shared Tree and Source Tree (Shortest-Path Tree)

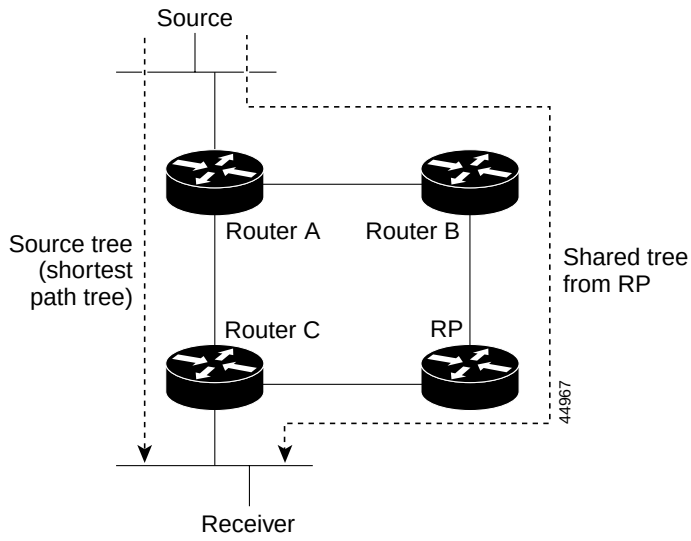

If the data rate warrants, leaf routers (routers without any downstream connections) on the shared tree can use the data distribution tree rooted at the source. This type of distribution tree is called a shortest-path tree or source tree. By default, the IOS software switches to a source tree upon receiving the first data packet from a source.

This process describes the move from a shared tree to a source tree:

1. A receiver joins a group; leaf Router C sends a join message toward the RP.
2. The RP puts a link to Router C in its outgoing interface list.
3. A source sends data; Router A encapsulates the data in a register message and sends it to the RP.
4. The RP forwards the data down the shared tree to Router C and sends a join message toward the source. At this point, data might arrive twice at Router C, once encapsulated and once natively.
5. When data arrives natively (unencapsulated) at the RP, it sends a register-stop message to Router A.
6. By default, reception of the first data packet prompts Router C to send a join message toward the source.
7. When Router C receives data on (S,G), it sends a prune message for the source up the shared tree.
8. The RP deletes the link to Router C from the outgoing interface of (S,G). The RP triggers a prune message toward the source.

Join and prune messages are sent for sources and RPs. They are sent hop-by-hop and are processed by each PIM device along the path to the source or RP. Register and register-stop messages are not sent hop-by-hop. They are sent by the designated router that is directly connected to a source and are received by the RP for the group.

Multiple sources sending to groups use the shared tree.

You can configure the PIM device to stay on the shared tree. For more information, see the "Delaying the Use of PIM Shortest-Path Tree" section on page 24-29.

Delaying the Use of PIM Shortest-Path Tree

The change from shared to source tree happens when the first data packet arrives at the last-hop router (Router C in Figure 24-9). This change occurs because the **ip pim spt-threshold** interface configuration command controls that timing; its default setting is 0 kbps.

The shortest-path tree requires more memory than the shared tree but reduces delay. You might want to postpone its use. Instead of allowing the leaf router to immediately move to the shortest-path tree, you can specify that the traffic must first reach a threshold.

You can configure when a PIM leaf router should join the shortest-path tree for a specified group. If a source sends at a rate greater than or equal to the specified kbps rate, the multilayer switch triggers a PIM join message toward the source to construct a source tree (shortest-path tree). If the traffic rate from the source drops below the threshold value, the leaf router switches back to the shared tree and sends a prune message toward the source.

You can specify to which groups the shortest-path tree threshold applies by using a group list (a standard access list). If a value of 0 is specified or if the group list is not used, the threshold applies to all groups.

Beginning in privileged EXEC mode, follow these steps to configure a traffic rate threshold that must be reached before multicast routing is switched from the source tree to the shortest-path tree:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, specify the multicast group to which the threshold will apply. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface on the leaf router that connects to the source tree.
Step 4	**ip pim spt-threshold** {*kbps* | **infinity**} [**group-list** *access-list-number*]	Specify the threshold that must be reached before moving to shortest-path tree (spt). • For *kbps*, specify the traffic rate in kilobits per second. The default is 0 kbps. The range is 0 to 4294967. • Specify **infinity** if you want all sources for the specified group to use the shared tree, never switching to the source tree. • (Optional) For **group-list** *access-list-number*, specify the access list created in Step 2. If the value is 0 or if the group-list is not used, the threshold applies to all groups.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default threshold, use the **no ip pim spt-threshold** interface configuration command.

Modifying the PIM Router-Query Message Interval

PIM routers and multilayer switches send PIM router-query messages to determine which device will be the DR for each LAN segment (subnet). The DR is responsible for sending IGMP host-query messages to all hosts on the directly connected LAN.

With PIM DM operation, the DR has meaning only if IGMPv1 is in use. IGMPv1 does not have an IGMP querier election process, so the elected DR functions as the IGMP querier. With PIM SM operation, the DR is the device that is directly connected to the multicast source. It sends PIM register messages to notify the RP that multicast traffic from a source needs to be forwarded down the shared tree. In this case, the DR is the device with the highest IP address.

By default, multicast routers and multilayer switches send PIM router-query messages every 30 seconds.

Beginning in privileged EXEC mode, follow these steps to modify the router-query message interval:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip pim query-interval** *seconds*	Configure the frequency at which the multilayer switch sends PIM router-query messages. The default is 30 seconds. The range is 1 to 65535.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default interval, use the **no ip pim query-interval** [*seconds*] interface configuration command.

Configuring Optional IGMP Features

This section describes how to configure optional IGMP features. It contains this configuration information:

- Default IGMP Configuration, page 24-31
- Changing the IGMP Version, page 24-32
- Changing the IGMP Query Timeout for IGMPv2, page 24-32
- Changing the Maximum Query Response Time for IGMPv2, page 24-33
- Configuring the Multilayer Switch as a Member of a Group, page 24-34
- Controlling Access to IP Multicast Groups, page 24-35
- Modifying the IGMP Host-Query Message Interval, page 24-36
- Configuring the Multilayer Switch as a Statically Connected Member, page 24-36

Default IGMP Configuration

Table 24-2 shows the default IGMP configuration.

Table 24-2 Default IGMP Configuration

Feature	Default Setting
IGMP version	Version 2 on all interfaces.
IGMP query timeout	60 seconds on all interfaces.
IGMP maximum query response time	10 seconds on all interfaces.
Multilayer switch as a member of a multicast group	No group memberships are defined.

Table 24-2 Default IGMP Configuration (continued)

Feature	Default Setting
Access to multicast groups	All groups are allowed on an interface.
IGMP host-query message interval	60 seconds on all interfaces.
Multilayer switch as a statically connected member	Disabled.

Changing the IGMP Version

By default, the multilayer switch uses IGMP Version 2, which allows features such as the IGMP query timeout and the maximum query response time.

All systems on the subnet must support the same version. The switch does not automatically detect Version 1 systems and switch to Version 1.

Configure the switch for Version 1 if your hosts do not support Version 2.

Beginning in privileged EXEC mode, follow these steps to change the IGMP version:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp version** {**2** | **1**}	Specify the IGMP version that the switch uses. **Note** If you change to version 1, you cannot configure the **ip igmp query-interval** or the **ip igmp query-max-response-time** interface configuration commands.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default version (2), use the **no ip igmp version** interface configuration command.

Changing the IGMP Query Timeout for IGMPv2

If you are using IGMPv2, you can specify the period of time before the multilayer switch takes over as the querier for the interface. By default, the switch waits twice the query interval controlled by the **ip igmp query-interval** interface configuration command. After that time, if the switch has received no queries, it becomes the querier.

You can determine the query interval by entering the **show ip igmp interface** *interface-id* privileged EXEC command.

Beginning in privileged EXEC mode, follow these steps to change the IGMP query timeout:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp querier-timeout** *seconds*	Specify the IGMP query timeout. The default is 60 seconds (twice the query interval). The range is 60 to 300.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default timeout value, use the **no ip igmp query-timeout** interface configuration command.

Changing the Maximum Query Response Time for IGMPv2

If you are using IGMPv2, you can change the maximum query response time advertised in IGMP queries. The maximum query response time allows the multilayer switch to quickly detect that there are no more directly connected group members on a LAN. Decreasing the value allows the switch to prune groups faster.

Beginning in privileged EXEC mode, follow these steps to change the maximum query response time:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp query-max-response-time** *seconds*	Change the maximum query response time advertised in IGMP queries. The default is 10 seconds. The range is 1 to 25.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default query-response time, use the **no ip igmp query-max-response-time** interface configuration command.

Configuring the Multilayer Switch as a Member of a Group

Multilayer switches can be configured as members of a multicast group. This is useful to determine multicast reachability in a network. If all the multicast-capable routers and multilayer switches that you administer are members of a multicast group, pinging that group causes all these devices to respond. The devices respond to ICMP echo-request packets addressed to a group of which they are members. Another example is the multicast trace-route tools provided in the Cisco IOS software.

Beginning in privileged EXEC mode, follow these steps to configure the multilayer switch to be a member of a group:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp join-group** *group-address*	Configure the switch to join a multicast group. By default, no group memberships are defined. For *group-address*, specify the multicast IP address in dotted decimal notation.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To cancel membership in a group, use the **no ip igmp join-group** *group-address* interface configuration command.

This example shows how to allow the switch to join multicast group 255.2.2.2:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip igmp join-group 255.2.2.2
```

Controlling Access to IP Multicast Groups

The multilayer switch sends IGMP host-query messages to determine which multicast groups have members on attached local networks. The switch then forwards to these group members all packets addressed to the multicast group. You can place a filter on each interface to restrict the multicast groups that hosts on the subnet serviced by the interface can join.

Beginning in privileged EXEC mode, follow these steps to filter multicast groups allowed on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp access-group** *access-list-number*	Specify the multicast groups that hosts on the subnet serviced by an interface can join. By default, all groups are allowed on an interface. For *access-list-number*, specify an IP standard access list number. The range is 1 to 99.
Step 4	**exit**	Return to global configuration mode.
Step 5	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list. • For *access-list-number*, specify the access list created in Step 3. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, specify the multicast group that hosts on the subnet can join. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 8	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable groups on an interface, use the **no ip igmp access-group** *access-list-number* interface configuration command.

This example shows how to configure hosts attached to Gigabit Ethernet interface 0/1 as able to join only group 255.2.2.2:

```
Switch(config)# access-list 1 255.2.2.2 0.0.0.0
Switch(config-if)# interface gigabitethernet0/1
Switch(config-if)# ip igmp access-group 1
```

Modifying the IGMP Host-Query Message Interval

The multilayer switch periodically sends IGMP host-query messages to discover which multicast groups are present on attached networks. These messages are sent to the all-hosts multicast group (224.0.0.1) with a time-to-live (TTL) of 1. The switch sends host-query messages to refresh its knowledge of memberships present on the network. If, after some number of queries, the IOS software discovers that no local hosts are members of a multicast group, the software stops forwarding multicast packets to the local network from remote origins for that group and sends a prune message upstream toward the source.

The switch elects a PIM designated router (DR) for the LAN (subnet). The DR is the router or multilayer switch with the highest IP address for IGMPv2; for IGMPv1, the DR is elected according to the multicast routing protocol that runs on the LAN. The designated router is responsible for sending IGMP host-query messages to all hosts on the LAN. In sparse mode, the designated router also sends PIM register and PIM join messages toward the RP router.

Beginning in privileged EXEC mode, follow these steps to modify the host-query interval:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp query-interval** *seconds*	Configure the frequency at which the designated router sends IGMP host-query messages. By default, the designated router sends IGMP host-query messages every 60 seconds to keep the IGMP overhead very low on hosts and networks. The range is 0 to 65535.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default frequency, use the **no ip igmp query-interval** interface configuration command.

Configuring the Multilayer Switch as a Statically Connected Member

Sometimes there is either no group member on a network segment or a host cannot report its group membership by using IGMP. However, you might want multicast traffic to go to that network segment. These are ways to pull multicast traffic down to a network segment:

- Use the **ip igmp join-group** interface configuration command. With this method, the multilayer switch accepts the multicast packets in addition to forwarding them. Accepting the multicast packets prevents the switch from fast switching.
- Use the **ip igmp static-group** interface configuration command. With this method, the multilayer switch does not accept the packets itself, but only forwards them. This method allows fast switching. The outgoing interface appears in the IGMP cache, but the switch itself is not a member, as evidenced by lack of an *L* (local) flag in the multicast route entry.

Beginning in privileged EXEC mode, follow these steps to configure the switch itself to be a statically connected member of a group (and allow fast switching):

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip igmp static-group** *group-address*	Configure the switch as a statically connected member of a group. By default, this feature is disabled.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show ip igmp interface** [*interface-id*]	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the switch as a member of the group, use the **no ip igmp static-group** interface configuration command.

Configuring Optional Multicast Routing Features

This section describes how to configure optional multicast routing features, which are grouped as follows:

- Features for Layer 2 connectivity and MBONE multimedia conference session and set up:
  - Enabling CGMP Server Support, page 24-38
  - Configuring sdr Listener Support, page 24-39
- Features that control bandwidth utilization:
  - Configuring the TTL Threshold, page 24-40
  - Configuring an IP Multicast Boundary, page 24-42

Enabling CGMP Server Support

The multilayer switch serves as a CGMP server for devices that do not support IGMP snooping but have CGMP client functionality. CGMP is a protocol used on Cisco routers and multilayer switches connected to Layer 2 Catalyst switches to perform tasks similar to those performed by IGMP. CGMP is necessary because the Layer 2 switch cannot distinguish between IP multicast data packets and IGMP report messages, which are both at the MAC-level and are addressed to the same group address.

Beginning in privileged EXEC mode, follow these steps to enable the CGMP server on the multilayer switch interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface that is connected to the Layer 2 Catalyst switch.
Step 3	**ip cgmp** [**proxy**]	Enable CGMP on the interface. By default, CGMP is disabled on all interfaces. Enabling CGMP triggers a CGMP join message. Enable CGMP only on Layer 3 interfaces connected to Layer 2 Catalyst switches. (Optional) When you enter the **proxy** keyword, the CGMP proxy function is enabled. The proxy router advertises the existence of non-CGMP-capable routers by sending a CGMP join message with the non-CGMP-capable router MAC address and a group address of 0000.0000.0000. **Note** To perform CGMP proxy, the multilayer switch must be the IGMP querier. If you configure the **ip cgmp proxy** command, you must manipulate the IP addresses so that the switch is the IGMP querier, which might be the highest or lowest IP address, depending on which version of IGMP is running on the network. An IGMP Version 2 querier is selected based on the lowest IP address on the interface. An IGMP Version 1 querier is selected based on the multicast routing protocol used on the interface.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.
Step 7		Verify the Layer 2 Catalyst switch CGMP-client configuration. For more information, refer to the documentation that shipped with the product.

To disable CGMP on the interface, use the **no ip cgmp** interface configuration command.

When multiple Cisco CGMP-capable devices are connected to a switched network and the **ip cgmp proxy** command is needed, we recommend that all devices be configured with the same CGMP option and have precedence for becoming the IGMP querier over non-Cisco routers.

Configuring sdr Listener Support

The MBONE is the small subset of Internet routers and hosts that are interconnected and capable of forwarding IP multicast traffic. Other interesting multimedia content is often broadcast over the MBONE. Before you can join a multimedia session, you need to know what multicast group address and port are being used for the session, when the session is going to be active, and what sort of applications (audio, video, and so forth) are required on your workstation. The MBONE Session Directory version 2 (sdr) tool provides this information. This freeware application can be downloaded from several sites on the World Wide Web, one of which is http://www.video.ja.net/mice/index.html.

SDR is a multicast application that listens to a well-known multicast group address and port for Session Announcement Protocol (SAP) multicast packets from SAP clients, which announce their conference sessions. These SAP packets contain a session description, the time the session is active, its IP multicast group addresses, media format, contact person, and other information about the advertised multimedia session. The information in the SAP packet is displayed in the SDR Session Announcement window.

Enabling sdr Listener Support

By default, the multilayer switch does not listen to session directory advertisements.

Beginning in privileged EXEC mode, follow these steps to enable the switch to join the default session directory group (224.2.127.254) on the interface and listen to session directory advertisements:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be enabled for sdr.
Step 3	**ip sdr listen**	Enable sdr listener support.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable sdr support, use the **no ip sdr listen** interface configuration command.

Limiting How Long an sdr Cache Entry Exists

By default, entries are never deleted from the sdr cache. You can limit how long the entry remains active so that if a source stops advertising SAP information, old advertisements are not needlessly kept.

Beginning in privileged EXEC mode, follow these steps to limit how long an sdr cache entry stays active in the cache:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip sdr cache-timeout** *minutes*	Limit how long an sdr cache entry stays active in the cache. By default, entries are never deleted from the cache. For *minutes*, specify a number from 1 to 4294967295.

	Command	Purpose
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip sdr cache-timeout** global configuration command. To delete the entire cache, use the **clear ip sdr** privileged EXEC command.

To display the session directory cache, use the **show ip sdr** privileged EXEC command.

Configuring the TTL Threshold

Each time an IP multicast packet is forwarded by the multilayer switch, the time-to-live (TTL) value in the IP header is decremented by one. If the packet TTL decrements to zero, the switch drops the packet. TTL thresholds can be applied to individual interfaces of the multilayer switch to prevent multicast packets with a TTL less than the TTL threshold from being forwarded out the interface. TTL thresholds provide a simple method to prevent the forwarding of multicast traffic beyond the boundary of a site or region, based on the TTL field in a multicast packet. This is known as TTL scoping.

Figure 24-10 shows a multicast packet arriving on Gigabit Ethernet interface 0/2 with a TTL value of 24. Assuming that the RPF check succeeds and that Gigabit Ethernet interfaces 0/1, 0/3, and 0/4 are all in the outgoing interface list, the packet would normally be forwarded out these interfaces. Because some TTL thresholds have been applied to these interfaces, the multilayer switch makes sure that the packet TTL value, which is decremented by 1 to 23, is greater than or equal to the interface TTL threshold before forwarding the packet out the interface. In this example, the packet is forwarded out interfaces 0/1 and 0/4, but not interface 0/3.

Figure 24-10 TTL Thresholds

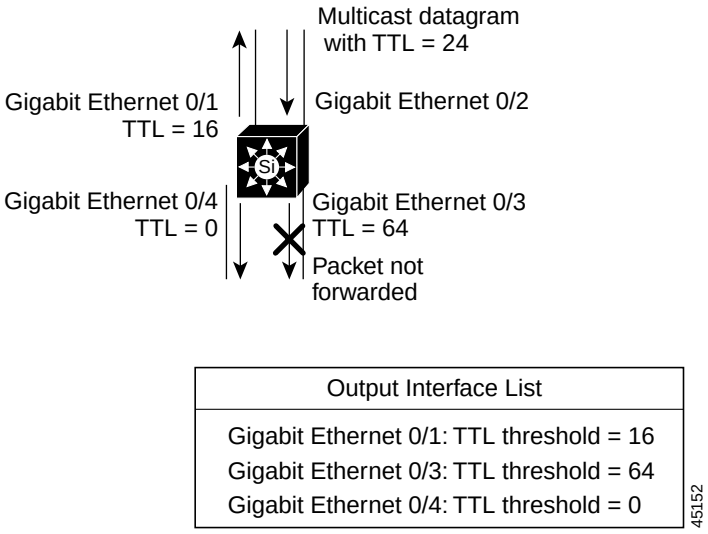

Figure 24-11 shows an example of TTL threshold boundaries being used to limit the forwarding of multicast traffic. Company XYZ has set a TTL threshold of 100 on all routed interfaces at the perimeter of its network. Multicast applications that constrain traffic to within the company's network need to send

multicast packets with an initial TTL value set to 99. The engineering and marketing departments have set a TTL threshold of 40 at the perimeter of their networks; therefore, multicast applications running on these networks can prevent their multicast transmissions from leaving their respective networks.

Figure 24-11 TTL Boundaries

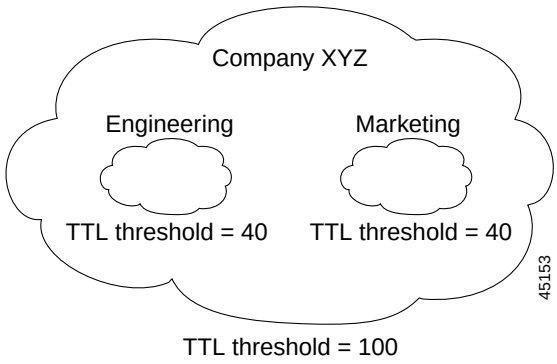

Beginning in privileged EXEC mode, follow these steps to change the default TTL threshold value:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip multicast ttl-threshold** *ttl-value*	Configure the TTL threshold of packets being forwarded out an interface. The default TTL value is 0 hops, which means that all multicast packets are forwarded out the interface. The range is 0 to 255. Only multicast packets with a TTL value greater than the threshold are forwarded out the interface. You should configure the TTL threshold only on routed interfaces at the perimeter of the network.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default TTL setting, use the **no ip multicast ttl-threshold** interface configuration command.

Configuring an IP Multicast Boundary

Like TTL thresholds, administratively-scoped boundaries can also be used to limit the forwarding of multicast traffic outside of a domain or subdomain. This approach uses a special range of multicast addresses, called *administratively-scoped addresses*, as the boundary mechanism. If you configure an administratively-scoped boundary on a routed interface, multicast traffic whose multicast group addresses fall in this range can not enter or exit this interface, thereby providing a firewall for multicast traffic in this address range.

Figure 24-12 shows that Company XYZ has an administratively-scoped boundary set for the multicast address range 239.0.0.0/8 on all routed interfaces at the perimeter of its network. This boundary prevents any multicast traffic in the range 239.0.0.0 through 239.255.255.255 from entering or leaving the network. Similarly, the engineering and marketing departments have an administratively-scoped boundary of 239.128.0.0/16 around the perimeter of their networks. This boundary prevents multicast traffic in the range of 239.128.0.0 through 239.128.255.255 from entering or leaving their respective networks.

Figure 24-12 Administratively-Scoped Boundaries

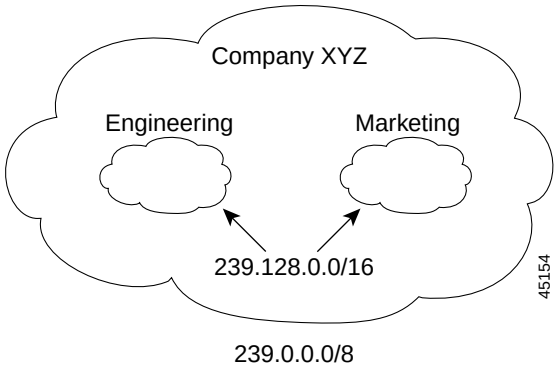

You can define an administratively-scoped boundary on a routed interface for multicast group addresses. A standard access list defines the range of addresses affected. When a boundary is defined, no multicast data packets are allowed to flow across the boundary from either direction. The boundary allows the same multicast group address to be reused in different administrative domains.

The IANA has designated the multicast address range 239.0.0.0 to 239.255.255.255 as the administratively-scoped addresses. This range of addresses can then be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

Beginning in privileged EXEC mode, follow these steps to set up an administratively-scoped boundary:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the number of the network or host from which the packet is being sent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 4	**ip multicast boundary** *access-list-number*	Configure the boundary, specifying the access list you created in Step 2.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the boundary, use the **no ip multicast boundary** interface configuration command.

This example shows how to set up a boundary for all administratively-scoped addresses:

```
Switch(config)# access-list 1 deny 239.0.0.0 0.255.255.255
Switch(config)# access-list 1 permit 224.0.0.0 15.255.255.255
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip multicast boundary 1
```

Configuring Basic DVMRP Interoperability Features

This section describes how to perform basic configuration tasks on your multilayer switch to interoperate with DVMRP devices. It contains this configuration information:

- Configuring DVMRP Interoperability, page 24-44
- Controlling Unicast Route Advertisements, page 24-44
- Configuring a DVMRP Tunnel, page 24-46
- Advertising Network 0.0.0.0 to DVMRP Neighbors, page 24-48
- Responding to mrinfo Requests, page 24-49

For more advanced DVMRP features, see the "Configuring Advanced DVMRP Interoperability Features" section on page 24-50.

Configuring DVMRP Interoperability

Cisco multicast routers and multilayer switches using PIM can interoperate with non-Cisco multicast routers that use the DVMRP.

PIM devices dynamically discover DVMRP multicast routers on attached networks by listening to DVMR probe messages. When a DVMRP neighbor has been discovered, the PIM device periodically sends DVMRP report messages advertising the unicast sources reachable in the PIM domain. By default, directly connected subnets and networks are advertised. The device forwards multicast packets that have been forwarded by DVMRP routers and, in turn, forwards multicast packets to DVMRP routers.

DVMRP interoperability is automatically activated when a Cisco PIM device receives a DVMRP probe message on a multicast-enabled interface. No specific IOS command is configured to enable DVMRP interoperability; however, you must enable multicast routing. For more information, see the "Configuring Basic Multicast Routing" section on page 24-15.

Controlling Unicast Route Advertisements

You should configure an access list on the PIM routed interface connected to the MBONE to limit the number of unicast routes that are advertised in DVMRP route reports; otherwise, all routes in the unicast routing table are advertised.

Note The mrouted protocol is a public-domain implementation of DVMRP. You must use mrouted Version 3.8 (which implements a nonpruning version of DVMRP) when Cisco routers and multilayer switches are directly connected to DVMRP routers or interoperate with DVMRP routers over an MBONE tunnel. DVMRP advertisements produced by the Cisco IOS software can cause older versions of the mrouted protocol to corrupt their routing tables and those of their neighbors.

You can configure what sources are advertised and what metrics are used by configuring the **ip dvmrp metric** interface configuration command. You can also direct all sources learned through a particular unicast routing process to be advertised into DVMRP.

Beginning in privileged EXEC mode, follow these steps to configure the sources that are advertised and the metrics that are used when DVMRP route-report messages are sent:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the number of the network or host from which the packet is being sent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface connected to the MBONE and enabled for multicast routing.
Step 4	**ip dvmrp metric** *metric* [**list** *access-list-number*] [[*protocol process-id*] | [**dvmrp**]]	Configure the metric associated with a set of destinations for DVMRP reports. • For *metric*, the range is 0 to 32. A value of 0 means that the route is not advertised. A value of 32 is equivalent to infinity (unreachable). • (Optional) For **list** *access-list-number*, enter the access list number created in Step 2. If specified, only the multicast destinations that match the access list are reported with the configured metric. • (Optional) For *protocol process-id*, enter the name of the unicast routing protocol, such as **eigrp**, **igrp**, **ospf**, **rip**, **static**, or **dvmrp**, and the process ID number of the routing protocol. If specified, only routes learned by the specified routing protocol are advertised in DVMRP report messages. • (Optional) If specified, the **dvmrp** keyword allows routes from the DVMRP routing table to be advertised with the configured *metric* or filtered.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the metric or route map, use the **no ip dvmrp metric** *metric* [**list** *access-list-number*] [[*protocol process-id*] | [**dvmrp**]] or the **no ip dvmrp metric** *metric* **route-map** *map-name* interface configuration command.

A more sophisticated way to achieve the same results as the preceding command is to use a route map (**ip dvmrp metric** *metric* **route-map** *map-name* interface configuration command) instead of an access list. You subject unicast routes to route-map conditions before they are injected into DVMRP.

This example shows how to configure DVMRP interoperability when the PIM device and the DVMRP router are on the same network segment. In this example, access list 1 advertises the networks (198.92.35.0, 198.92.36.0, 198.92.37.0, 131.108.0.0, and 150.136.0.0) to the DVMRP router, and access list 2 prevents all other networks from being advertised (**ip dvmrp metric 0** interface configuration command).

```
Switch(config-if)# interface gigabitethernet0/1
Switch(config-if)# ip address 131.119.244.244 255.255.255.0
Switch(config-if)# ip pim dense-mode
Switch(config-if)# ip dvmrp metric 1 list 1
Switch(config-if)# ip dvmrp metric 0 list 2
Switch(config-if)# exit
Switch(config)# access-list 1 permit 198.92.35.0 0.0.0.255
Switch(config)# access-list 1 permit 198.92.36.0 0.0.0.255
Switch(config)# access-list 1 permit 198.92.37.0 0.0.0.255
Switch(config)# access-list 1 permit 131.108.0.0 0.0.255.255
Switch(config)# access-list 1 permit 150.136.0.0 0.0.255.255
Switch(config)# access-list 1 deny   0.0.0.0 255.255.255.255
Switch(config)# access-list 2 permit 0.0.0.0 255.255.255.255
```

Configuring a DVMRP Tunnel

The Cisco IOS software supports DVMRP tunnels to the MBONE. You can configure a DVMRP tunnel on a router or multilayer switch if the other end is running DVMRP. The software then sends and receives multicast packets through the tunnel. This strategy allows a PIM domain to connect to the DVMRP router when all routers on the path do not support multicast routing. You cannot configure a DVMRP tunnel between two routers.

When a Cisco router or multilayer switch runs DVMRP through a tunnel, it advertises sources in DVMRP report messages, much as it does on real networks. The software also caches DVMRP report messages it receives and uses them in its Reverse Path Forwarding (RPF) calculation. This behavior allows the software to forward multicast packets received through the tunnel.

When you configure a DVMRP tunnel, you should assign an IP address to a tunnel in these cases:

- To send IP packets through the tunnel
- To configure the Cisco IOS software to perform DVMRP summarization

The software does not advertise subnets through the tunnel if the tunnel has a different network number from the subnet. In this case, the software advertises only the network number through the tunnel.

Beginning in privileged EXEC mode, follow these steps to configure a DVMRP tunnel:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create a standard access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the number of the network or host from which the packet is being sent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 3	**interface tunnel** *number*	Enter interface configuration mode, and specify a tunnel interface.
Step 4	**tunnel source** *ip-address*	Specify the source address of the tunnel interface. Enter the IP address of the interface on the multilayer switch.
Step 5	**tunnel destination** *ip-address*	Specify the destination address of the tunnel interface. Enter the IP address of the mrouted router.
Step 6	**tunnel mode dvmrp**	Configure the encapsulation mode for the tunnel to DVMRP.
Step 7	**ip address** *address mask* or **ip unnumbered** *type number*	Assign an IP address to the interface. or Configure the interface as unnumbered.
Step 8	**ip pim** [**dense-mode** | **sparse-mode**]	Configure the PIM mode on the interface.
Step 9	**ip dvmrp accept-filter** *access-list-number* [*distance*] **neighbor-list** *access-list-number*	Configure an acceptance filter for incoming DVMRP reports. By default, all destination reports are accepted with a distance of 0. Reports from all neighbors are accepted. • For *access-list-number*, specify the access list number created in Step 2. Any sources that match the access list are stored in the DVMRP routing table with distance. • (Optional) For *distance*, enter the administrative distance to the destination. By default, the administrative distance for DVMRP routes is 0 and take precedence over unicast routing table routes. If you have two paths to a source, one through unicast routing (using PIM as the multicast routing protocol) and another using DVMRP, and if you want to use the PIM path, increase the administrative distance for DVMRP routes. The range is 1 to 255. • For **neighbor-list** *access-list-number*, enter the number of the neighbor list created in Step 2. DVMRP reports are accepted only by those neighbors on the list.
Step 10	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 11	**show running-config**	Verify your entries.
Step 12	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable the filter, use the **no ip dvmrp accept-filter** *access-list-number* [*distance*] **neighbor-list** *access-list-number* interface configuration command.

This example shows how to configure a DVMRP tunnel. In this configuration, the IP address of the tunnel on the Cisco multilayer switch is assigned *unnumbered*, which causes the tunnel to appear to have the same IP address as Gigabit Ethernet interface 0/1. The tunnel endpoint source address is 172.16.2.1, and the tunnel endpoint address of the remote DVMRP router to which the tunnel is connected is 192.168.1.10. Any packets sent through the tunnel are encapsulated in an outer IP header. The Cisco multilayer switch is configured to accept incoming DVMRP reports with a distance of 100 from 198.92.37.0 through 198.92.37.255.

```
Switch(config)# ip multicast-routing
Switch(config)# interface tunnel 0
Switch(config-if)# ip unnumbered gigabitethernet 0/1
Switch(config-if)# ip pim dense-mode
Switch(config-if)# tunnel source gigabitethernet 0/1
Switch(config-if)# tunnel destination 192.168.1.10
Switch(config-if)# tunnel mode dvmrp
Switch(config-if)# ip dvmrp accept-filter 1 100
Switch(config-if)# interface gigabitethernet 0/1
Switch(config-if)# ip address 172.16.2.1 255.255.255.0
Switch(config-if)# ip pim dense-mode
Switch(config)# exit
Switch(config)# access-list 1 permit 198.92.37.0 0.0.0.255
```

Advertising Network 0.0.0.0 to DVMRP Neighbors

If your multilayer switch is a neighbor of an mrouted version 3.6 device, you can configure the Cisco IOS software to advertise network 0.0.0.0 (the default route) to the DVMRP neighbor. The DVMRP default route computes the RPF information for any multicast sources that do not match a more specific route.

Do not advertise the DVMRP default into the MBONE.

Beginning in privileged EXEC mode, follow these steps to advertise network 0.0.0.0 to DVMRP neighbors on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface that is connected to the DVMRP router.

	Command	Purpose
Step 3	**ip dvmrp default-information** {**originate** | **only**}	Advertise network 0.0.0.0 to DVMRP neighbors. Use this command only when the multilayer switch is a neighbor of mrouted version 3.6 machines. The keywords have these meanings: • **originate**—Specifies that other routes more specific than 0.0.0.0 can also be advertised. • **only**—Specifies that no DVMRP routes other than 0.0.0.0 are advertised.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To prevent the default route advertisement, use the **no ip dvmrp default-information** {**originate** | **only**} interface configuration command.

Responding to mrinfo Requests

The Cisco IOS software answers mrinfo requests sent by mrouted systems and Cisco routers and multilayer switches. The software returns information about neighbors through DVMRP tunnels and all the routed interfaces. This information includes the metric (always set to 1), the configured TTL threshold, the status of the interface, and various flags. You can also use the **mrinfo** privileged EXEC command to query the router or switch itself, as in this example:

```
Switch# mrinfo
  171.69.214.27 (mm1-7kd.cisco.com) [version cisco 11.1] [flags: PMS]:
  171.69.214.27 -> 171.69.214.26 (mm1-r7kb.cisco.com) [1/0/pim/querier]
  171.69.214.27 -> 171.69.214.25 (mm1-45a.cisco.com) [1/0/pim/querier]
  171.69.214.33 -> 171.69.214.34 (mm1-45c.cisco.com) [1/0/pim]
  171.69.214.137 -> 0.0.0.0 [1/0/pim/querier/down/leaf]
  171.69.214.203 -> 0.0.0.0 [1/0/pim/querier/down/leaf]
  171.69.214.18 -> 171.69.214.20 (mm1-45e.cisco.com) [1/0/pim]
  171.69.214.18 -> 171.69.214.19 (mm1-45c.cisco.com) [1/0/pim]
  171.69.214.18 -> 171.69.214.17 (mm1-45a.cisco.com) [1/0/pim]
```

Configuring Advanced DVMRP Interoperability Features

Cisco routers and multilayer switches run PIM to forward multicast packets to receivers and receive multicast packets from senders. It is also possible to propagate DVMRP routes into and through a PIM cloud. PIM uses this information; however, Cisco routers and multilayer switches do not implement DVMRP to forward multicast packets.

This section describes how to perform advanced optional configuration tasks on your multilayer switch to interoperate with DVMRP devices. It contains this configuration information:

- Enabling DVMRP Unicast Routing, page 24-50
- Rejecting a DVMRP Nonpruning Neighbor, page 24-51
- Controlling Route Exchanges, page 24-53

For information on basic DVMRP features, see the "Configuring Basic DVMRP Interoperability Features" section on page 24-43.

Enabling DVMRP Unicast Routing

Because multicast routing and unicast routing require separate topologies, PIM must follow the multicast topology to build loopless distribution trees. Using DVMRP unicast routing, Cisco routers, multilayer switches, and mrouted-based machines exchange DVMRP unicast routes, to which PIM can then reverse-path forward.

Cisco devices do not perform DVMRP multicast routing among each other, but they can exchange DVMRP routes. The DVMRP routes provide a multicast topology that might differ from the unicast topology. This allows PIM to run over the multicast topology, thereby allowing sparse-mode PIM over the MBONE topology.

When DVMRP unicast routing is enabled, the router or switch caches routes learned in DVMRP report messages in a DVMRP routing table. When PIM is running, these routes might be preferred over routes in the unicast routing table, allowing PIM to run on the MBONE topology when it is different from the unicast topology.

DVMRP unicast routing can run on all interfaces. For DVMRP tunnels, it uses DVMRP multicast routing. This feature does not enable DVMRP multicast routing among Cisco routers and multilayer switches. However, if there is a DVMRP-capable multicast router, the Cisco device can do PIM/DVMRP multicast routing.

Beginning in privileged EXEC mode, follow these steps to enable DVMRP unicast routing:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface that is connected to the DVMRP router.
Step 3	**ip dvmrp unicast-routing**	Enable DVMRP unicast routing (to send and receive DVMRP routes). This feature is disabled by default.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable this feature, use the **no ip dvmrp unicast-routing** interface configuration command.

Rejecting a DVMRP Nonpruning Neighbor

By default, Cisco devices accept all DVMRP neighbors as peers, regardless of their DVMRP capability. However, some non-Cisco devices run old versions of DVMRP that cannot prune, so they continuously receive forwarded packets, wasting bandwidth. Figure 24-13 shows this scenario.

Figure 24-13 Leaf Nonpruning DVMRP Neighbor

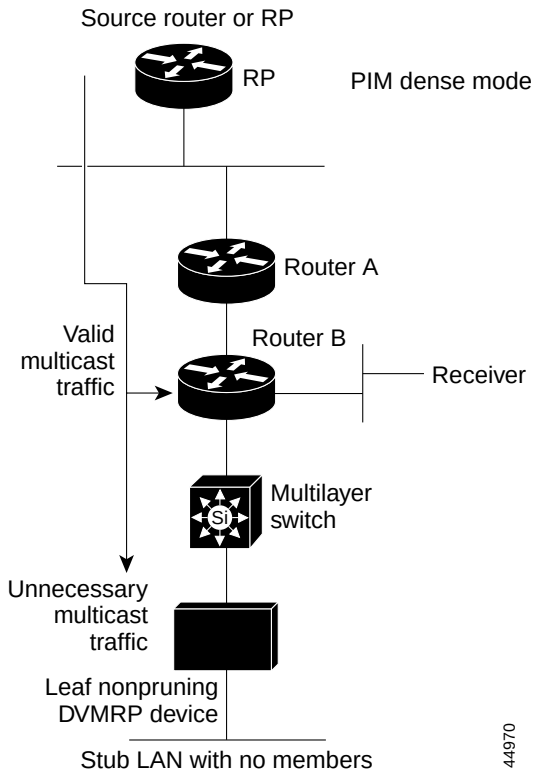

You can prevent the multilayer switch from peering (communicating) with a DVMRP neighbor if that neighbor does not support DVMRP pruning or grafting. To do so, configure the multilayer switch (which is a neighbor to the leaf, nonpruning DVMRP machine) with the **ip dvmrp reject-non-pruners** interface configuration command on the interface connected to the nonpruning machine as shown in Figure 24-14. In this case, when the multilayer switch receives DVMRP probe or report message without the prune-capable flag set, the switch logs a syslog message and discards the message.

Figure 24-14 Router Rejects Nonpruning DVMRP Neighbor

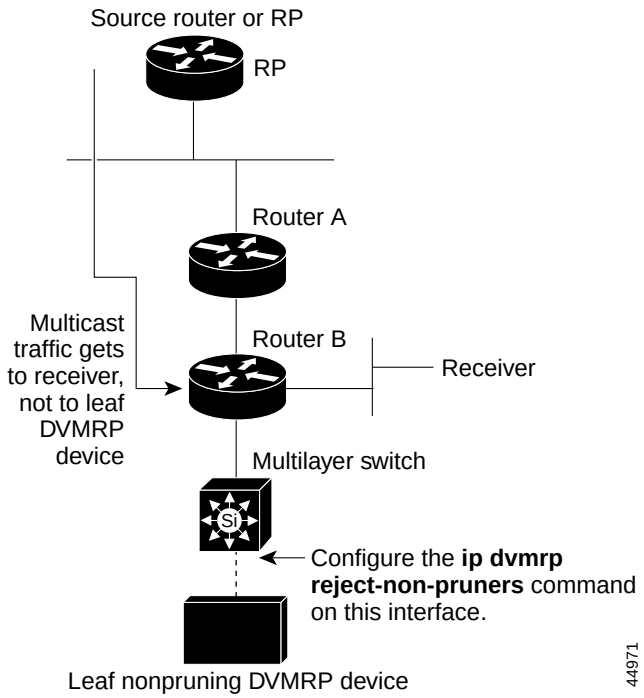

Note that the **ip dvmrp reject-non-pruners** interface configuration command prevents peering with neighbors only. If there are any nonpruning routers multiple hops away (downstream toward potential receivers) that are not rejected, a nonpruning DVMRP network might still exist.

Beginning in privileged EXEC mode, follow these steps to prevent peering with nonpruning DVMRP neighbors:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface connected to the nonpruning DVMRP neighbor.
Step 3	**ip dvmrp reject-non-pruners**	Prevent peering with nonpruning DVMRP neighbors.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To disable this function, use the **no ip dvmrp reject-non-pruners** interface configuration command.

Controlling Route Exchanges

This section describes how to tune the Cisco device advertisements of DVMRP routes. It contains this configuration information:

- Limiting the Number of DVMRP Routes Advertised, page 24-53
- Changing the DVMRP Route Threshold, page 24-54
- Configuring a DVMRP Summary Address, page 24-54
- Disabling DVMRP Autosummarization, page 24-56
- Adding a Metric Offset to the DVMRP Route, page 24-56

Limiting the Number of DVMRP Routes Advertised

By default, only 7000 DVMRP routes are advertised over an interface enabled to run DVMRP (that is, a DVMRP tunnel, an interface where a DVMRP neighbor has been discovered, or an interface configured to run the **ip dvmrp unicast-routing** interface configuration command).

Beginning in privileged EXEC mode, follow these steps to change the DVMRP route limit:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip dvmrp route-limit** *count*	Change the number of DVMRP routes advertised over an interface enabled for DVMRP. This command prevents misconfigured **ip dvmrp metric** interface configuration commands from causing massive route injection into the MBONE. By default, 7000 routes are advertised. The range is 0 to 4294967295.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To configure no route limit, use the **no ip dvmrp route-limit** global configuration command.

Changing the DVMRP Route Threshold

By default, 10,000 DVMRP routes can be received per interface within a 1-minute interval. When that rate is exceeded, a syslog message is issued, warning that there might be a route surge occurring. The warning is typically used to quickly detect when devices have been misconfigured to inject a large number of routes into the MBONE.

Beginning in privileged EXEC mode, follow these steps to change the threshold number of routes that trigger the warning:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip dvmrp routehog-notification** *route-count*	Configure the number of routes that trigger a syslog message. The default is 10,000 routes. The range is 1 to 4294967295.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default route count, use the **no ip dvmrp routehog-notification** global configuration command.

Use the **show ip igmp interface** privileged EXEC command to display a running count of routes. When the count is exceeded, *** *ALERT* *** is appended to the line.

Configuring a DVMRP Summary Address

By default, a Cisco device advertises in DVMRP route-report messages only connected unicast routes (that is, only routes to subnets that are directly connected to the router) from its unicast routing table. These routes undergo normal DVMRP classful route summarization. This process depends on whether the route being advertised is in the same classful network as the interface over which it is being advertised.

Figure 24-15 shows an example of the default behavior. This example shows that the DVMRP report sent by the Cisco router contains the three original routes received from the DVMRP router that have been poison-reversed by adding 32 to the DVMRP metric. Listed after these routes are two routes that are advertisements for the two directly connected networks (176.32.10.0/24 and 176.32.15.0/24) that were taken from the unicast routing table. Because the DVMRP tunnel shares the same IP address as Fast Ethernet 0/1 and falls into the same Class B network as the two directly connected subnets, classful summarization of these routes was not performed. As a result, the DVMRP router is able to poison-reverse only these two routes to the directly connected subnets and is able to only RPF properly for multicast traffic sent by sources on these two Ethernet segments. Any other multicast source in the network behind the Cisco router that is not on these two Ethernet segments does not properly RPF-check on the DVMRP router and is discarded.

You can force the Cisco router to advertise the summary address (specified by the address and mask pair in the **ip dvmrp summary-address** *address mask* interface configuration command) in place of any route that falls in this address range. The summary address is sent in a DVMRP route report if the unicast routing table contains at least one route in this range; otherwise, the summary address is not advertised. In Figure 24-15, you configure the **ip dvmrp summary-address** command on the Cisco router tunnel interface. As a result, the Cisco router sends only a single summarized Class B advertisement for network 176.32.0.0.16 from the unicast routing table.

Figure 24-15 Only Connected Unicast Routes Are Advertised by Default

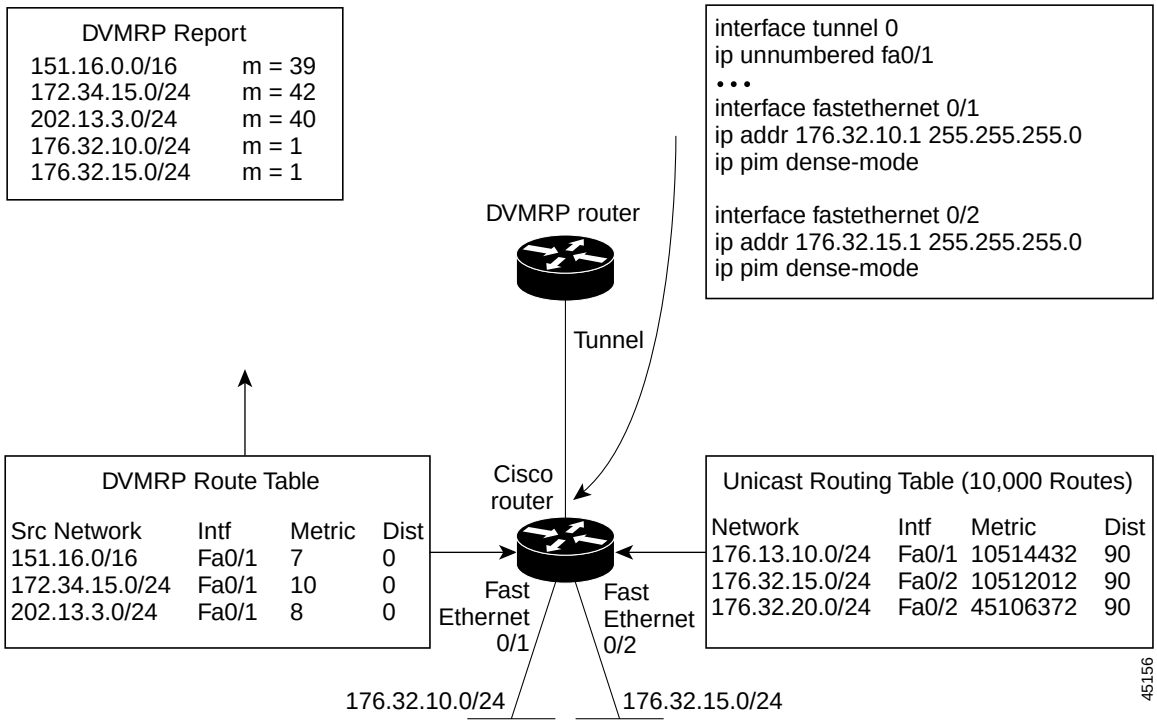

Beginning in privileged EXEC mode, follow these step to customize the summarization of DVMRP routes if the default classful autosummarization does not suit your needs:

Note At least one more-specific route must be present in the unicast routing table before a configured summary address is advertised.

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration command, and specify the interface that is connected to the DVMRP router.
Step 3	**ip dvmrp summary-address** *address mask* [**metric** *value*]	Specify a DVMRP summary address. • For **summary-address** *address mask*, specify the summary IP address and mask that is advertised instead of the more specific route. • (Optional) For **metric** *value*, specify the metric that is advertised with the summary address. The default is 1. The range is 1 to 32.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the summary address, use the **no ip dvmrp summary-address** *address mask* [**metric** *value*] interface configuration command.

Disabling DVMRP Autosummarization

By default, the Cisco IOS software automatically performs some level of DVMRP summarization. Disable this function if you want to advertise all routes, not just a summary. In some special cases, you can use the neighboring DVMRP router with all subnet information to better control the flow of multicast traffic in the DVMRP network. One such case might occur if the PIM network is connected to the DVMRP cloud at several points and more specific (unsummarized) routes are being injected into the DVMRP network to advertise better paths to individual subnets inside the PIM cloud.

If you configure the **ip dvmrp summary-address** interface configuration command and did not configure **no ip dvmrp auto-summary**, you get both custom and autosummaries.

Beginning in privileged EXEC mode, follow these steps to disable DVMRP autosummarization:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface connected to the DVMRP router.
Step 3	**no ip dvmrp auto-summary**	Disable DVMRP autosummarization.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To re-enable auto summarization, use the **ip dvmrp auto-summary** interface configuration command.

Adding a Metric Offset to the DVMRP Route

By default, the multilayer switch increments by 1 the metric (hop count) of a DVMRP route advertised in incoming DVMRP reports. You can change the metric if you want to favor or not favor a certain route.

For example, a route is learned by multilayer switch A, and the same route is learned by multilayer switch B with a higher metric. If you want to use the path through switch B because it is a faster path, you can apply a metric offset to the route learned by switch A to make it larger than the metric learned by switch B, and you can choose the path through switch B.

Beginning in privileged EXEC mode, follow these steps to change the default metric:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to be configured.
Step 3	**ip dvmrp metric-offset** [**in** | **out**] *increment*	Change the metric added to DVMRP routes advertised in incoming reports. The keywords have these meanings: • (Optional) **in**—Specifies that the increment value is added to incoming DVMRP reports and is reported in mrinfo replies. • (Optional) **out**—Specifies that the increment value is added to outgoing DVMRP reports for routes from the DVMRP routing table. If neither **in** nor **out** is specified, **in** is the default. For *increment*, specify the value that is added to the metric of a DVMRP router advertised in a report message. The range is 1 to 31. If the **ip dvmrp metric-offset** command is not configured on an interface, the default increment value for incoming routes is 1, and the default for outgoing routes is 0.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no ip dvmrp metric-offset** interface configuration command.

Monitoring and Maintaining IP Multicast Routing

This section describes how to monitor and maintain IP multicast routing. It contains this configuration information:

- Clearing Caches, Tables, and Databases, page 24-58
- Displaying System and Network Statistics, page 24-58
- Monitoring IP Multicast Routing, page 24-59

Clearing Caches, Tables, and Databases

You can remove all contents of a particular cache, table, or database. Clearing a cache, table, or database might be necessary when the contents of the particular structure are or suspected to be invalid.

You can use any of the privileged EXEC commands in Table 24-3 to clear IP multicast caches, tables, and databases:

Table 24-3 Commands for Clearing Caches, Tables, and Databases

Command	Purpose
clear ip cgmp	Clear all group entries the Catalyst switches have cached.
clear ip dvmrp route {***** | *route*}	Delete routes from the DVMRP routing table.
clear ip igmp group [*group-name* | *group-address* | *interface*]	Delete entries from the IGMP cache.
clear ip mroute {***** | *group* [*source*]}	Delete entries from the IP multicast routing table.
clear ip pim auto-rp *rp-address*	Clear the Auto-RP cache.
clear ip sdr [*group-address* | **"***session-name***"**]	Delete the Session Directory Protocol Version 2 cache or an sdr cache entry.

Displaying System and Network Statistics

You can display specific statistics such as the contents of IP routing tables, caches, and databases. You can display information to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path your device's packets are taking through the network.

You can use any of the privileged EXEC commands in Table 24-4 to display various routing statistics:

Table 24-4 Commands for Displaying System and Network Statistics

Command	Purpose
ping [*group-name* | *group-address*]	Send an ICMP Echo Request to a multicast group address.
show ip dvmrp route [*ip-address*]	Display the entries in the DVMRP routing table.
show ip igmp groups [*group-name* | *group-address* | *type number*]	Display the multicast groups that are directly connected to the multilayer switch and that were learned through IGMP.
show ip igmp interface [*type number*]	Display multicast-related information about an interface.
show ip mcache [*group* [*source*]]	Display the contents of the IP fast-switching cache.switching, displaying
show ip mpacket [*source-address* | *name*] [*group-address* | *name*] [**detail**]	Display the contents of the circular cache-header buffer.
show ip mroute [*group-name* | *group-address*] [*source*] [**summary**] [**count**] [**active** *kbps*]	Display the contents of the IP multicast routing table.

Table 24-4 Commands for Displaying System and Network Statistics (continued)

Command	Purpose
show ip pim interface [*type number*] [**count**]	Display information about interfaces configured for PIM.
show ip pim neighbor [*type number*]	List the PIM neighbors discovered by the multilayer switch.
show ip pim rp [*group-name* | *group-address*]	Display the RP routers associated with a sparse-mode multicast group.
show ip rpf {*source-address* | *name*}	Display how the multilayer switch is doing Reverse-Path Forwarding (that is, from the unicast routing table, DVMRP routing table, or static mroutes).
show ip sdr [*group* | **"***session-name***"** | **detail**]	Display the Session Directory Protocol Version 2 cache.

Monitoring IP Multicast Routing

You can use the privileged EXEC commands in Table 24-5 to monitor IP multicast routers, packets, and paths:

Table 24-5 Commands for Monitoring IP Multicast Routing

Command	Purpose
mrinfo [*hostname* | *address*] [*source-address* | *interface*]	Query a multicast router or multilayer switch about which neighboring multicast devices are peering with it.
mstat *source* [*destination*] [*group*]	Display IP multicast packet rate and loss information.
mtrace *source* [*destination*] [*group*]	Trace the path from a source to a destination branch for a multicast distribution tree for a given group.

CHAPTER 25

Configuring MSDP

This chapter describes how to configure the Multicast Source Discovery Protocol (MSDP) on your multilayer switch. The MSDP connects multiple Protocol-Independent Multicast sparse-mode (PIM-SM) domains.

MSDP is not fully supported in this IOS release because of a lack of support for Multicast Border Gateway Protocol (MBGP), which works closely with MSDP. However, it is possible to create default peers that MSDP can operate with if MBGP is not running.

To use this feature, you must have the enhanced multilayer software (EMI) image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding MSDP, page 25-1
- Configuring MSDP, page 25-4
- Monitoring and Maintaining MSDP, page 25-19

Understanding MSDP

MSDP allows multicast sources for a group to be known to all rendezvous points (RPs) in different domains. Each PIM-SM domain uses its own RPs and does not depend on RPs in other domains. An RP runs MSDP over the Transmission Control Protocol (TCP) to discover multicast sources in other domains.

An RP in a PIM-SM domain has an MSDP peering relationship with MSDP-enabled devices in another domain. The peering relationship occurs over a TCP connection, primarily exchanging a list of sources sending to multicast groups. The TCP connections between RPs are achieved by the underlying routing system. The receiving RP uses the source lists to establish a source path.

The purpose of this topology is to have domains discover multicast sources in other domains. If the multicast sources are of interest to a domain that has receivers, multicast data is delivered over the normal, source-tree building mechanism in PIM-SM. MSDP is also used to announce sources sending to a group. These announcements must originate at the domain RP.

MSDP depends heavily on the Border Gateway Protocol (BGP) or MBGP for interdomain operation. We recommend that you run MSDP in RPs in your domain that are RPs for sources sending to global groups to be announced to the Internet.

MSDP Operation

Figure 25-1 shows MSDP operating between two MSDP peers. PIM uses MSDP as the standard mechanism to register a source with the RP of a domain. When MSDP is configured, this sequence occurs.

When a source sends its first multicast packet, the first-hop router (*designated router* or RP) directly connected to the source sends a PIM register message to the RP. The RP uses the register message to register the active source and to forward the multicast packet down the shared tree in the local domain. With MSDP configured, the RP also forwards a source-active (SA) message to all MSDP peers. The SA message identifies the source, the group the source is sending to, and the address of the RP or the originator ID (the IP address of the interface used as the RP address), if configured.

Each MSDP peer receives and forwards the SA message away from the originating RP to achieve peer-RPF flooding. The MSDP device examines the BGP or MBGP routing table to determine which peer is the next hop toward the originating RP of the SA message. Such a peer is called an *RPF peer* (reverse-path forwarding peer). The MSDP device forwards the message to all MSDP peers other than the RPF peer. For information on how to configure an MSDP peer when BGP and MBGP are not supported, see the "Configuring a Default MSDP Peer" section on page 25-4.

If the MSDP peer receives the same SA message from a non-RPF peer toward the originating RP, it drops the message. Otherwise, it forwards the message to all its MSDP peers.

When the RP for a domain receives the SA message from an MSDP peer, it determines if it has any join requests for the group the SA message describes. If the (*,G) entry exists with a nonempty outgoing interface list, the domain is interested in the group, and the RP triggers an (S,G) join toward the source. After the (S,G) join reaches the source's DR, a branch of the source tree has been built from the source to the RP in the remote domain. Multicast traffic can now flow from the source across the source tree to the RP and then down the shared tree in the remote domain to the receiver.

Figure 25-1 MSDP Running Between RP Peers

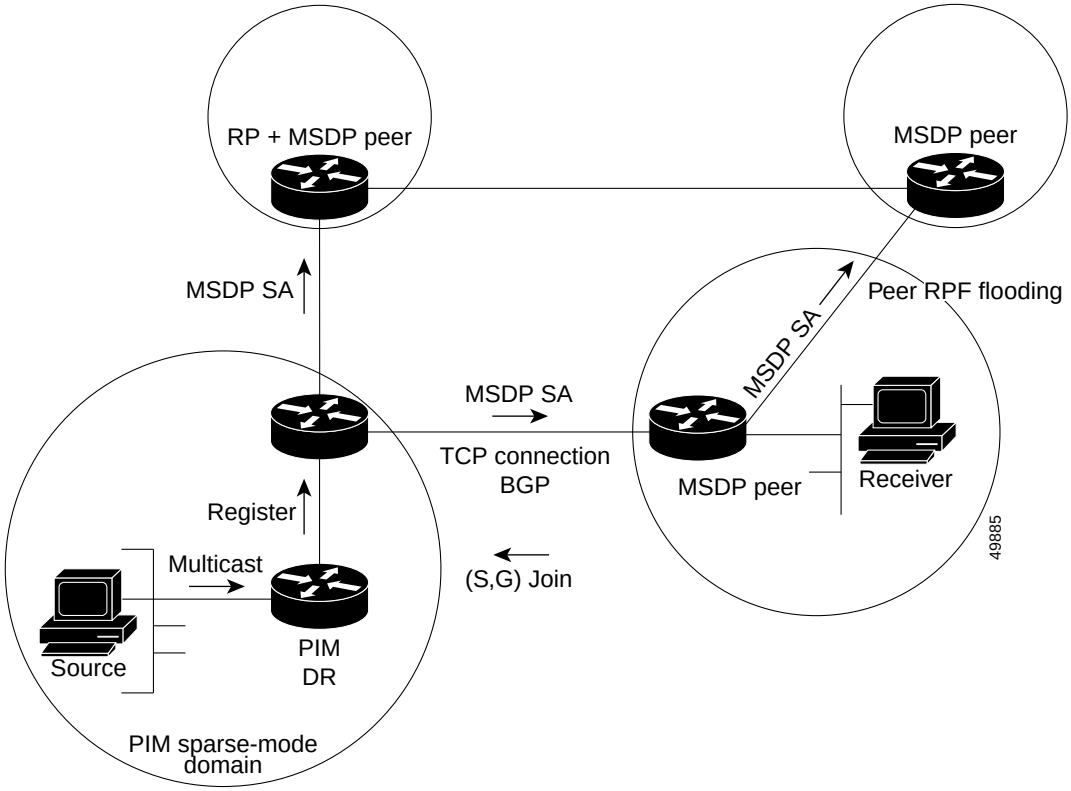

MSDP Benefits

MSDP has these benefits:

- It breaks up the shared multicast distribution tree. You can make the shared tree local to your domain. Your local members join the local tree, and join messages for the shared tree never need to leave your domain.
- PIM sparse-mode domains can rely only on their own RPs, decreasing reliance on RPs in another domain. This increases security because you can prevent your sources from being known outside your domain.
- Domains with only receivers can receive data without globally advertising group membership.
- Global source multicast routing table state is not required, saving memory.

Configuring MSDP

This section describes how to configure MSDP. It contains this configuration information:

- Default MSDP Configuration, page 25-4
- Configuring a Default MSDP Peer, page 25-4 (required)
- Caching Source-Active State, page 25-6 (optional)
- Requesting Source Information from an MSDP Peer, page 25-8 (optional)
- Controlling Source Information that Your Switch Originates, page 25-8 (optional)
- Controlling Source Information that Your Switch Forwards, page 25-12 (optional)
- Controlling Source Information that Your Switch Receives, page 25-14 (optional)
- Configuring an MSDP Mesh Group, page 25-16 (optional)
- Shutting Down an MSDP Peer, page 25-16 (optional)
- Including a Bordering PIM Dense-Mode Region in MSDP, page 25-17 (optional)
- Configuring an Originating Address other than the RP Address, page 25-18 (optional)

Default MSDP Configuration

MSDP is not enabled, and no default MSDP peer exists.

Configuring a Default MSDP Peer

In this IOS release, because BGP and MBGP are not supported, you cannot configure an MSDP peer on the local multilayer switch by using the **ip msdp peer** global configuration command. Instead, you define a default MSDP peer (by using the **ip msdp default-peer** global configuration command) from which to accept all SA messages for the multilayer switch. The default MSDP peer must be a previously configured MSDP peer. Configure a default MSDP peer when the multilayer switch is not BGP- or MBGP-peering with an MSDP peer. If a single MSDP peer is configured, the multilayer switch always accepts all SA messages from that peer.

Figure 25-2 shows a network in which default MSDP peers might be used. In Figure 25-2, a customer who owns Multilayer Switch B is connected to the Internet through two Internet service providers (ISPs), one owning Router A and the other owning Router C. They are not running BGP or MBGP between them. To learn about sources in the ISP's domain or in other domains, multilayer Switch B at the customer site identifies Router A as its default MSDP peer. Multilayer Switch B advertises SA messages to both Router A and Router C but accepts SA messages only from Router A or only Router C. If Router A is first in the configuration file, it is used if it is running. If Router A is not running, only then does multilayer Switch B accept SA messages from Router C. This is the default behavior without a prefix list.

If you specify a prefix list, the peer is a default peer only for the prefixes in the list. You can have multiple active default peers when you have a prefix list associated with each. When you do not have any prefix lists, you can configure multiple default peers, but only the first one is the active default peer as long as the router has connectivity to this peer and the peer is alive. If the first configured peer fails or the connectivity to this peer fails, the second configured peer becomes the active default, and so on.

The ISP probably uses a prefix list to define which prefixes it accepts from the customer's router.

Figure 25-2 Default MSDP Peer Network

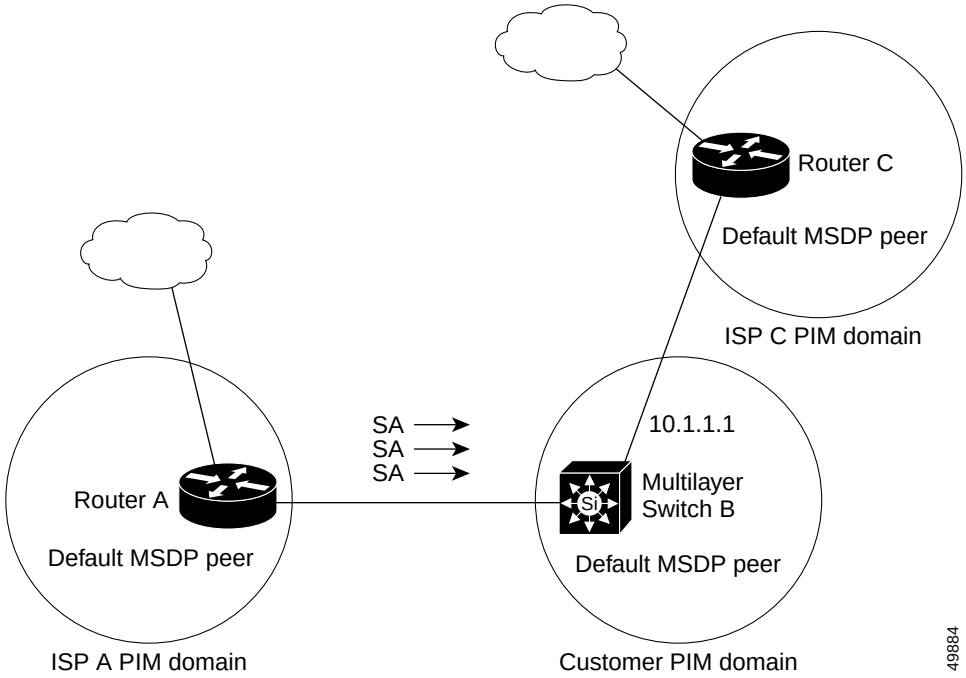

Beginning in privileged EXEC mode, follow these steps to specify a default MSDP peer:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp default-peer** *ip-address* | *name* [**prefix-list** *list*]	Define a default peer from which to accept all MSDP SA messages. • For *ip-address* | *name*, enter the IP address or Domain Name System (DNS) server name of the MSDP default peer. • (Optional) For **prefix-list** *list*, enter the list name that specifies the peer to be the default peer only for the listed prefixes. You can have multiple active default peers when you have a prefix list associated with each. When you enter multiple **ip msdp default-peer** commands with the **prefix-list** keyword, you use all the default peers at the same time for different RP prefixes. This syntax is typically used in a service provider cloud that connects stub site clouds. When you enter multiple **ip msdp default-peer** commands without the **prefix-list** keyword, a single active peer accepts all SA messages. If that peer fails, the next configured default peer accepts all SA messages. This syntax is typically used at a stub site.

	Command	Purpose
Step 3	**ip prefix-list** *name* [**description** *string*] | **seq** *number* {**permit** | **deny**} *network length*	(Optional) Create a prefix list using the name specified in Step 2. • (Optional) For **description** *string*, enter a description of up to 80 characters to describe this prefix list. • For **seq** *number*, enter the sequence number of the entry. The range is 1 to 4294967294. • The **deny** keyword denies access to matching conditions. • The **permit** keyword permits access to matching conditions. • For *network length*, specify the network number and length (in bits) of the network mask that is permitted or denied.
Step 4	**ip msdp description** {*peer-name* | *peer-address*} *text*	(Optional) Configure a description for the specified peer to make it easier to identify in a configuration or in **show** command output. By default, no description is associated with an MSDP peer.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the default peer, use the **no ip msdp default-peer** global configuration command.

This example shows a partial configuration of Router A and Router C in Figure 25-2. Each of these ISPs have more than one customer (like the customer in Figure 25-2) who use default peering (no BGP or MBGP). In that case, they might have similar configurations. That is, they accept SAs only from a default peer if the SA is permitted by the corresponding prefix list.

Router A

```
Router(config)# ip msdp default-peer 10.1.1.1
Router(config)# ip msdp default-peer 10.1.1.1 prefix-list site-a
Router(config)# ip prefix-list site-b permit 10.0.0.0/8
```

Router C

```
Router(config)# ip msdp default-peer 10.1.1.1 prefix-list site-a
Router(config)# ip prefix-list site-b permit 10.0.0.0/8
```

Caching Source-Active State

By default, the multilayer switch does not cache source/group pairs from received SA messages. When the switch forwards the MSDP SA information, it does not store it in memory. Therefore, if a member joins a group soon after a SA message is received by the local RP, that member needs to wait until the next SA message to hear about the source. This delay is known as join latency.

If you want to sacrifice some memory in exchange for reducing the latency of the source information, you can configure the switch to cache SA messages.

Beginning in privileged EXEC mode, follow these steps to enable the caching of source/group pairs:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp cache-sa-state** [**list** *access-list-number*]	Enable the caching of source/group pairs (create an SA state). Those pairs that pass the access list are cached. For **list** *access-list-number*, the range is 100 to 199.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	Create an IP extended access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 100 to 199. Enter the same number created in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *protocol*, enter **ip** as the protocol name. • For *source*, enter the number of the network or host from which the packet is being sent. • For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. • For *destination*, enter the number of the network or host to which the packet is being sent. • For *destination-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the destination. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note An alternative to this command is the **ip msdp sa-request** global configuration command, which causes the multilayer switch to send an SA request message to the MSDP peer when a new member for a group becomes active. For more information, see the next section.

To return to the default setting (no SA state is created), use the **no ip msdp cache-sa-state** global configuration command.

This example shows how to enable the cache state for all sources in 171.69.0.0/16 sending to groups 224.2.0.0/16:

```
Switch(config)# ip msdp cache-sa-state 100
Switch(config)# access-list 100 permit ip 171.69.0.0 0.0.255.255 224.2.0.0 0.0.255.255
```

Requesting Source Information from an MSDP Peer

Local RPs can send SA requests and get immediate responses for all active sources for a given group. By default, the multilayer switch does not send any SA request messages to its MSDP peers when a new member joins a group and wants to receive multicast traffic. The new member waits to receive the next periodic SA message.

If you want a new member of a group to learn the current, active multicast sources in a connected PIM sparse-mode domain that are sending to a group, configure the switch to send SA request messages to the specified MSDP peer when a new member joins a group. The peer replies with the information in its SA cache. If the peer does not have a cache configured, this command has no result. Configuring this feature reduces join latency but sacrifices memory.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send SA request messages to the MSDP peer when a new member joins a group and wants to receive multicast traffic:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp sa-request** {*ip-address* | *name*}	Configure the switch to send SA request messages to the specified MSDP peer. For *ip-address* | *name*, enter the IP address or name of the MSDP peer from which the local switch requests SA messages when a new member for a group becomes active. Repeat the command for each MSDP peer that you want to supply with SA messages.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp sa-request** {*ip-address* | *name*} global configuration command.

This example shows how to configure the switch to send SA request messages to the MSDP peer at 171.69.1.1:

```
Switch(config)# ip msdp sa-request 171.69.1.1
```

Controlling Source Information that Your Switch Originates

You can control the multicast source information that originates with your switch:

- Sources you advertise (based on your sources)
- Receivers of source information (based on knowing the requestor)

For more information, see the "Redistributing Sources" section on page 25-9 and the "Filtering Source-Active Request Messages" section on page 25-11.

Redistributing Sources

SA messages are originated on RPs to which sources have registered. By default, any source that registers with an RP is advertised. The *A flag* is set in the RP when a source is registered, which means the source is advertised in an SA unless it is filtered.

Beginning in privileged EXEC mode, follow these steps to further restrict which registered sources are advertised:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp redistribute** [**list** *access-list-name*] [**asn** *aspath-access-list-number*] [**route-map** *map*]	Configure which (S,G) entries from the multicast routing table are advertised in SA messages. By default, only sources within the local domain are advertised. • (Optional) For **list** *access-list-name*, enter the name or number of an IP standard or extended access list. The range is 1 to 99 for standard access lists and 100 to 199 for extended lists. The access list controls which local sources are advertised and to which groups they send. • (Optional) For **asn** *aspath-access-list-number*, enter the IP standard or extended access list number in the range 1 to 199. This access list number must also be configured in the **ip as-path access-list** command. • (Optional) For **route-map** *map*, enter the IP standard or extended access list number in the range 1 to 199. This access list number must also be configured in the **ip as-path access-list** command. The access list or autonomous system path access list determines which (S,G) pairs are advertised.

	Command	Purpose
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create an IP standard access list, repeating the command as many times as necessary.
	or	or
	access-list *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	Create an IP extended access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99 for standard access lists and 100 to 199 for extended lists. Enter the same number created in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *protocol*, enter **ip** as the protocol name. • For *source*, enter the number of the network or host from which the packet is being sent. • For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. • For *destination*, enter the number of the network or host to which the packet is being sent. • For *destination-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the destination. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp redistribute** global configuration command.

Filtering Source-Active Request Messages

By default, only multilayer switches that are caching SA information can respond to SA requests. By default, such a switch honors all SA request messages from its MSDP peers and supplies the IP addresses of the active sources.

However, you can configure the switch to ignore all SA requests from an MSDP peer. You can also honor only those SA request messages from a peer for groups described by a standard access list. If the groups in the access list pass, SA request messages are accepted. All other such messages from the peer for other groups are ignored.

Beginning in privileged EXEC mode, follow these steps to configure one of these options:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp filter-sa-request** *ip-address* | *name* or **ip msdp filter-sa-request** {*ip-address* | *name*} **list** *access-list-number*	Filter all SA request messages from the specified MSDP peer. or Filter SA request messages from the specified MSDP peer for groups that pass the standard access list. The access list describes a multicast group address. The range for the access-list-number is 1 to 99.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *source* [*source-wildcard*]	Create an IP standard access list, repeating the command as many times as necessary. • For *access-list-number*, the range is 1 to 99. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *source*, enter the number of the network or host from which the packet is being sent. • (Optional) For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp filter-sa-request** {*ip-address* | *name*} global configuration command.

This example shows how to configure the switch to filter SA request messages from the MSDP peer at 171.69.2.2. SA request messages from sources on network 192.4.22.0 pass access list 1 and are accepted; all others are ignored.

```
Switch(config)# ip msdp filter sa-request 171.69.2.2 list 1
Switch(config)# access-list 1 permit 192.4.22.0 0.0.0.255
```

Controlling Source Information that Your Switch Forwards

By default, the multilayer switch forwards all SA messages it receives to all its MSDP peers. However, you can prevent outgoing messages from being forwarded to a peer by using a filter or by setting a time-to-live (TTL) value. These methods are described in the next sections.

Using a Filter

By creating a filter, you can perform one of these actions:

- Filter all source/group pairs
- Specify an IP extended access list to pass only certain source/group pairs
- Filter based on match criteria in a route map

Beginning in privileged EXEC mode, follow these steps to apply a filter:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp sa-filter out** *ip-address* | *name*	Filter all SA messages to the specified MSDP peer.
	or	or
	ip msdp sa-filter out {*ip-address* | *name*} **list** *access-list-number*	To the specified peer, pass only those SA messages that pass the IP extended access list. The range for the extended *access-list-number* is 100 to 199. If both the **list** and the **route-map** keywords are used, all conditions must be true to pass any (S,G) pair in outgoing SA messages.
	or	or
	ip msdp sa-filter out {*ip-address* | *name*} **route-map** *map-tag*	To the specified MSDP peer, pass only those SA messages that meet the match criteria in the route map *map-tag*. If all match criteria are true, a **permit** from the route map passes routes through the filter. A **deny** filters routes.

	Command	Purpose
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	(Optional) Create an IP extended access list, repeating the command as many times as necessary. • For *access-list-number*, enter the number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *protocol*, enter **ip** as the protocol name. • For *source*, enter the number of the network or host from which the packet is being sent. • For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. • For *destination*, enter the number of the network or host to which the packet is being sent. • For *destination-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the destination. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp sa-filter out** {*ip-address* | *name*} [**list** *access-list-number*] [**route-map** *map-tag*] global configuration command.

This example shows how to allow only (S,G) pairs that pass access list 100 to be forwarded in an SA message to the peer named *switch.cisco.com*:

```
Switch(config)# ip msdp peer switch.cisco.com connect-source gigabitethernet0/1
Switch(config)# ip msdp sa-filter out switch.cisco.com list 100
Switch(config)# access-list 100 permit ip 171.69.0.0 0.0.255.255 224.20 0 0.0.255.255
```

Using TTL to Limit the Multicast Data Sent in SA Messages

You can use a TTL value to control what data is encapsulated in the first SA message for every source. Only multicast packets with an IP-header TTL greater than or equal to the *ttl* argument are sent to the specified MSDP peer. For example, you can limit internal traffic to a TTL of 8. If you want other groups to go to external locations, you must send those packets with a TTL greater than 8.

Beginning in privileged EXEC mode, follow these steps to establish a TTL threshold:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp ttl-threshold** {*ip-address* | *name*} *ttl*	Limit which multicast data is encapsulated in the first SA message to the specified MSDP peer. • For *ip-address* | *name*, enter the IP address or name of the MSDP peer to which the TTL limitation applies. • For *ttl*, enter the TTL value. The default is 0, which means all multicast data packets are forwarded to the peer until the TTL is exhausted. The range is 0 to 255.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no ip msdp ttl-threshold** {*ip-address* | *name*} global configuration command.

Controlling Source Information that Your Switch Receives

By default, the multilayer switch receives all SA messages that its MSDP Reverse-Path Forwarding peers send to it. However, you can control the source information that you receive from MSDP peers by filtering incoming SA messages. In other words, you can configure the switch to not accept them.

You can perform one of these actions:

- Filter all incoming SA messages from an MSDP peer
- Specify an IP extended access list to pass certain source/group pairs
- Filter based on match criteria in a route map

Beginning in privileged EXEC mode, follow these steps to apply a filter:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp sa-filter in** *ip-address* | *name*	Filter all SA messages from the specified MSDP peer.
	or	or
	ip msdp sa-filter in {*ip-address* | *name*} **list** *access-list-number*	From the specified peer, pass only those SA messages that pass the IP extended access list. The range for the extended *access-list-number* is 100 to 199. If both the **list** and the **route-map** keywords are used, all conditions must be true to pass any (S,G) pair in incoming SA messages.
	or	or
	ip msdp sa-filter in {*ip-address* | *name*} **route-map** *map-tag*	From the specified MSDP peer, pass only those SA messages that meet the match criteria in the route map *map-tag*. If all match criteria are true, a **permit** from the route map passes routes through the filter. A **deny** will filter routes.
Step 3	**access-list** *access-list-number* {**deny** | **permit**} *protocol source source-wildcard destination destination-wildcard*	(Optional) Create an IP extended access list, repeating the command as many times as necessary. • For *access-list-number*, enter the number specified in Step 2. • The **deny** keyword denies access if the conditions are matched. The **permit** keyword permits access if the conditions are matched. • For *protocol*, enter **ip** as the protocol name. • For *source*, enter the number of the network or host from which the packet is being sent. • For *source-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. • For *destination*, enter the number of the network or host to which the packet is being sent. • For *destination-wildcard*, enter the wildcard bits in dotted decimal notation to be applied to the destination. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove the filter, use the **no ip msdp sa-filter in** {*ip-address* | *name*} [**list** *access-list-number*] [**route-map** *map-tag*] global configuration command.

This example shows how to filter all SA messages from the peer named *switch.cisco.com*:

```
Switch(config)# ip msdp peer switch.cisco.com connect-source gigabitethernet0/1
Switch(config)# ip msdp sa-filter in switch.cisco.com
```

Configuring an MSDP Mesh Group

An MSDP mesh group is a group of MSDP speakers that have fully meshed MSDP connectivity among one another. Any SA messages received from a peer in a mesh group are not forwarded to other peers in the same mesh group. Thus, you reduce SA message flooding and simplify peer-RPF flooding. Use the **ip msdp mesh-group** global configuration command when there are multiple RPs within a domain. It is especially used to send SA messages across a domain. You can configure multiple mesh groups (with different names) in a single multilayer switch.

Beginning in privileged EXEC mode, follow these steps to create a mesh group:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp mesh-group** *name* {*ip-address* | *name*}	Configure an MSDP mesh group, and specify the MSDP peer belonging to that mesh group. By default, the MSDP peers do not belong to a mesh group. • For *name*, enter the name of the mesh group. • For *ip-address* | *name*, enter the IP address or name of the MSDP peer to be a member of the mesh group.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.
Step 6		Repeat this procedure on each MSDP peer in the group.

To remove an MSDP peer from a mesh group, use the **no ip msdp mesh-group** *name* {*ip-address* | *name*} global configuration command.

Shutting Down an MSDP Peer

If you want to configure many MSDP commands for the same peer and you do not want the peer to become active, you can shut down the peer, configure it, and later bring it up. When a peer is shut down, the TCP connection is terminated and is not restarted. You can also shut down an MSDP session without losing configuration information for the peer.

Beginning in privileged EXEC mode, follow these steps to shut down a peer:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp shutdown** {*peer-name* | *peer address*}	Administratively shut down the specified MSDP peer without losing configuration information. For *peer-name* | *peer address*, enter the IP address or name of the MSDP peer to shut down.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To bring the peer back up, use the **no ip msdp shutdown** {*peer-name* | *peer address*} global configuration command. The TCP connection is reestablished

Including a Bordering PIM Dense-Mode Region in MSDP

You can configure MSDP on a multilayer switch that borders a PIM sparse-mode region with a dense-mode region. By default, active sources in the dense-mode region do not participate in MSDP.

Note We do not recommend using the **ip msdp border sa-address** global configuration command. It is better to configure the border router in the sparse-mode domain to proxy-register sources in the dense-mode domain to the RP of the sparse-mode domain and have the sparse-mode domain use standard MSDP procedures to advertise these sources.

Beginning in privileged EXEC mode, follow these steps to configure the border router to send SA messages for sources active in the dense-mode region to the MSDP peers:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp border sa-address** *type number*	Configure the switch on the border between a dense-mode and sparse-mode region to send SA messages about active sources in the dense-mode region. For *type number*, specify the interface type and number from which the IP address is derived and used as the RP address in SA messages. The IP address of the interface is used as the Originator-ID, which is the RP field in the SA message.
Step 3	**ip msdp redistribute** [**list** *access-list-name*] [**asn** *aspath-access-list-number*] [**route-map** *map*]	Configure which (S,G) entries from the multicast routing table are advertised in SA messages. For more information, see the "Redistributing Sources" section on page 25-9.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show running-config**	Verify your entries.
Step 6	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

Note that the **ip msdp originator-id** global configuration command also identifies an interface type and number to be used as the RP address. If both the **ip msdp border sa-address** and the **ip msdp originator-id** global configuration commands are configured, the address derived from the **ip msdp originator-id** command determines the RP address.

To return to the default setting (active sources in the dense-mode region do not participate in MSDP), use the **no ip msdp border sa-address** *type number* global configuration command.

Configuring an Originating Address other than the RP Address

You can allow an MSDP speaker that originates an SA message to use the IP address of the interface as the RP address in the SA message by changing the Originator ID. You might change the Originator ID in one of these cases:

- If you configure a logical RP on multiple multilayer switches in an MSDP mesh group.
- If you have a multilayer switch that borders a PIM sparse-mode domain and a dense-mode domain. If a switch borders a dense-mode domain for a site, and sparse-mode is being used externally, you might want dense-mode sources to be known to the outside world. Because this switch is not an RP, it would not have an RP address to use in an SA message. Therefore, this command provides the RP address by specifying the address of the interface.

Beginning in privileged EXEC mode, follow these steps to allow an MSDP speaker that originates an SA message to use the IP address on the interface as the RP address in the SA message:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**ip msdp originator-id** *type number*	Configures the RP address in SA messages to be the address of the originating device interface. For *type number*, specify the interface type and number on the local switch.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entries.
Step 5	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

If both the **ip msdp border sa-address** and the **ip msdp originator-id** global configuration commands are configured, the address derived from the **ip msdp originator-id** command determines the address of the RP.

To prevent the RP address from being derived in this way, use the **no ip msdp originator-id** *type number* global configuration command.

Monitoring and Maintaining MSDP

To monitor MSDP SA messages, peers, state, or peer status, use one or more of the privileged EXEC commands in Table 25-1:

Table 25-1 Commands for Monitoring and Maintaining MSDP

Command	Purpose
debug ip msdp [*peer-address* | *name*] [**detail**] [**routes**]	Debugs an MSDP activity.
debug ip msdp resets	Debugs MSDP peer reset reasons.
show ip msdp count [*autonomous-system-number*]	Displays the number of sources and groups originated in SA messages from each autonomous system. The **ip msdp cache-sa-state** command must be configured for this command to produce any output.
show ip msdp peer [*peer-address* | *name*]	Displays detailed information about an MSDP peer.
show ip msdp sa-cache [*group-address* | *source-address* | *group-name* | *source-name*] [*autonomous-system-number*]	Displays (S,G) state learned from MSDP peers.
show ip msdp summary	Displays MSDP peer status and SA message counts.

To clear MSDP connections, statistics, or SA cache entries, use the privileged EXEC commands in Table 25-2:

Table 25-2 Commands for Clearing MSDP Connections, Statistics, or SA Cache Entries

Command	Purpose
clear ip msdp peer *peer-address* | *name*	Clears the TCP connection to the specified MSDP peer, resetting all MSDP message counters.
clear ip msdp statistics [*peer-address* | *name*]	Clears statistics counters for one or all the MSDP peers without resetting the sessions.
clear ip msdp sa-cache [*group-address* | *name*]	Clears the SA cache entries for all entries, all sources for a specific group, or all entries for a specific source/group pair.

CHAPTER 26

Configuring Fallback Bridging

This chapter describes how to configure fallback bridging on your switch. With fallback bridging, you can forward non-IP protocols that the multilayer switch does not route between VLAN bridge domains and routed ports.

To use this feature, you must have the enhanced multilayer software (EMI) image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the EMI installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Bridging and IBM Networking Command Reference for Release 12.1*.

This chapter consists of these sections:

- Understanding Fallback Bridging, page 26-1
- Configuring Fallback Bridging, page 26-3
- Monitoring and Maintaining the Network, page 26-12

Understanding Fallback Bridging

With fallback bridging, the switch bridges together two or more VLANs or routed ports, essentially connecting multiple VLANs within one bridge domain. Fallback bridging forwards traffic that the multilayer switch does not route and forwards traffic belonging to a nonroutable protocol such as DECnet.

Fallback bridging does not allow the spanning trees from the VLANs being bridged to collapse; each VLAN has its own Spanning Tree Protocol (STP) instance and a separate spanning tree, called the VLAN-bridge spanning tree, which runs on top of the bridge group to prevent loops.

A VLAN bridge domain is represented using the switch virtual interface (SVI). A set of SVIs and routed ports (which do not have any VLANs associated with them) can be configured to form a bridge group.

Recall that an SVI represents a VLAN of switch ports as one interface to the routing or bridging function in the system. Only one SVI can be associated with a VLAN, and it is only necessary to configure an SVI for a VLAN when you want to route between VLANs, to fallback-bridge nonroutable protocols between VLANs, or to provide IP host connectivity to the switch. A routed port is a physical port that

acts like a port on a router, but it is not connected to a router. A routed port is not associated with a particular VLAN, does not support VLAN subinterfaces, but behaves like a normal routed interface. For more information about SVIs and routed ports, see Chapter 8, "Configuring Interface Characteristics."

A bridge group is an internal organization of network interfaces on a switch. Bridge groups cannot be used to identify traffic switched within the bridge group outside the switch on which they are defined. Bridge groups on the same switch function as distinct bridges; that is, bridged traffic and bridge protocol data units (BPDUs) cannot be exchanged between different bridge groups on a switch. An interface can be a member of only one bridge group. Use a bridge group for each separately bridged (topologically distinct) network connected to the switch.

The purpose of placing network interfaces into a bridge group is twofold:

- To bridge all nonrouted traffic among the network interfaces making up the bridge group. If the packet destination address is in the bridge table, it is forwarded on a single interface in the bridge group. If the packet destination address is not in the bridge table, it is flooded on all forwarding interfaces in the bridge group. The bridge places source addresses in the bridge table as it learns them during the bridging process.
- To participate in the spanning-tree algorithm by receiving, and in some cases sending, BPDUs on the LANs to which they are attached. A separate spanning process runs for each configured bridge group. Each bridge group participates in a separate spanning-tree instance. A bridge group establishes a spanning-tree instance based on the BPDUs it receives on only its member interfaces.

Figure 26-1 shows a fallback bridging network example. The multilayer switch has two interfaces configured as SVIs with different assigned IP addresses and attached to two different VLANs. Another interface is configured as a routed port with its own IP address. If all three of these ports are assigned to the same bridge group, non-IP protocol frames can be forwarded among the end stations connected to the switch.

Figure 26-1 Fallback Bridging Network Example

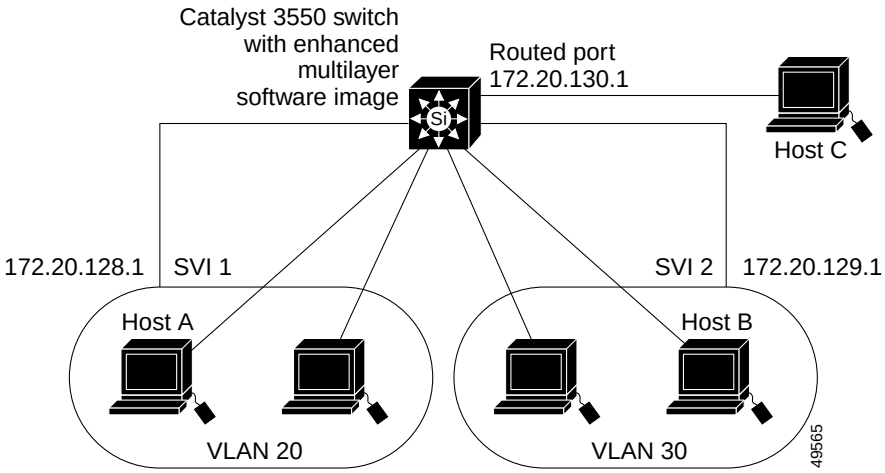

Configuring Fallback Bridging

This section describes how to configure fallback bridging on your switch. It contains this configuration information:

- Default Fallback Bridging Configuration, page 26-3
- Creating a Bridge Group, page 26-4
- Preventing the Forwarding of Dynamically Learned Stations, page 26-5
- Configuring the Bridge Table Aging Time, page 26-6
- Filtering Frames by a Specific MAC Address, page 26-6
- Adjusting Spanning-Tree Parameters, page 26-7

Default Fallback Bridging Configuration

Table 26-1 shows the default fallback bridging configuration.

Table 26-1 Default Fallback Bridging Configuration

Feature	Default Setting
Bridge groups	None are defined or assigned to an interface. No VLAN-bridge STP is defined.
Switch forwards frames for stations that it has dynamically learned	Enabled.
Bridge table aging time for dynamic entries	300 seconds.
MAC-layer frame filtering	Disabled.
Spanning tree parameters:	
• Switch priority	• 32768.
• Interface priority	• 128.
• Interface path cost	• 10 Mbps: 100. 100 Mbps: 19. 1000 Mbps: 4.
• Hello BPDU interval	• 2 seconds.
• Forward-delay interval	• 20 seconds.
• Maximum idle interval	• 30 seconds.

Creating a Bridge Group

To configure fallback bridging for a set of SVIs or routed ports, these interfaces must be assigned to bridge groups. All interfaces in the same group belong to the same bridge domain. Each SVI or routed port can be assigned to only one bridge group. A maximum of 31 bridge groups can be configured on the switch.

Note The protected port feature is not compatible with fallback bridging. When fallback bridging is enabled, it is possible for packets to be forwarded from one protected port on a switch to another protected port on the same switch if the ports are in different VLANs.

Beginning in privileged EXEC mode, follow these steps to create a bridge group and assign an interface to it:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **protocol vlan-bridge**	Assign a bridge group number, and specify the VLAN-bridge spanning-tree protocol to run in the bridge group. The **ibm** and **dec** keywords are not supported. For *bridge-group*, specify the bridge group number. The range is 1 to 255. You can create up to 31 bridge groups. Frames are bridged only among interfaces in the same group.
Step 3	**interface** *interface-id*	Enter interface configuration mode, and specify the interface on which you want to assign the bridge group. The specified interface must be one of these: • A routed port: a physical port that you have configured as a Layer 3 port by entering the **no switchport** interface configuration command. • An SVI: a VLAN interface that you created by using the **interface vlan** *vlan-id* global configuration command. These ports must have IP addresses assigned to them. For more information, see the "Configuring Layer 3 Interfaces" section on page 8-22.
Step 4	**bridge-group** *bridge-group*	Assign the interface to the bridge group created in Step 2. By default, the interface is not assigned to any bridge group. An interface can be assigned to only one bridge group.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show running-config**	Verify your entries.
Step 7	**copy running-config startup-config**	(Optional) Save your entries in the configuration file.

To remove a bridge group, use the **no bridge** *bridge-group* **protocol vlan-bridge** global configuration command. To remove an interface from a bridge group, use the **no bridge-group** *bridge-group* interface configuration command.

This example shows how to create bridge group 10, specify the VLAN-bridge STP to run in the bridge group, and assign an interface to the bridge group:

```
Switch(config)# bridge 10 protocol vlan-bridge
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no switchport
Switch(config-if)# bridge-group 10
```

Preventing the Forwarding of Dynamically Learned Stations

By default, the switch forwards any frames for stations that it has dynamically learned. By disabling this activity, the switch only forwards frames whose addresses have been statically configured into the forwarding cache.

Beginning in privileged EXEC mode, follow these steps to prevent the switch from forwarding frames for stations that it has dynamically learned:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**no bridge** *bridge-group* **acquire**	Enable the switch to stop forwarding any frames for stations that it has dynamically learned through the discovery process and to limit frame forwarding to statically configured stations. The switch filters all frames except those whose destined-to addresses have been statically configured into the forwarding cache. To configure a static address, use the **bridge** *bridge-group* **address** *mac-address* {**forward** | **discard**} global configuration command. For *bridge-group*, specify the bridge group number. The range is 1 to 255.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To cause the switch to forward frames to stations that it has dynamically learned, use the **bridge** *bridge-group* **acquire** global configuration command.

This example shows how to prevent the switch from forwarding frames for stations that it has dynamically learned in bridge group 10:

```
Switch(config)# no bridge 10 acquire
```

Configuring the Bridge Table Aging Time

A switch forwards, floods, or drops packets based on the bridge table. The bridge table maintains both static and dynamic entries. Static entries are entered by you or learned by the switch. Dynamic entries are entered by the bridge learning process. A dynamic entry is automatically removed after a specified length of time, known as aging time, from the time the entry was created or last updated.

If you are likely to move hosts on a switched network, decrease the aging-time to enable the switch to quickly adapt to the change. If hosts on a switched network do not continuously send packets, increase the aging time to keep the dynamic entries for a longer time and thus reduce the possibility of flooding when the hosts send again.

Beginning in privileged EXEC mode, follow these steps to configure the aging time:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **aging-time** *seconds*	Specify the length of time that a dynamic entry remains in the bridge table from the time the entry was created or last updated. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 0 to 1000000. The default is 300 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To return to the default aging-time interval, use the **no bridge** *bridge-group* **aging-time** global configuration command.

This example shows how to change the bridge table aging time to 200 seconds for bridge group 10:

```
Switch(config)# bridge 10 aging-time 200
```

Filtering Frames by a Specific MAC Address

A switch examines frames and sends them through the internetwork according to the destination address; a switch does not forward a frame back to its originating network segment. You can use the software to configure specific administrative filters that filter frames based on information other than the paths to their destinations.

You can filter frames with a particular MAC-layer station destination address. Any number of addresses can be configured in the system without a performance penalty.

Beginning in privileged EXEC mode, follow these steps to filter by the MAC-layer address:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **address** *mac-address* {**forward** | **discard**} [*interface-id*]	Specify the MAC address to discard or forward. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For **address** *mac-address*, specify the MAC-layer destination address to be filtered. • Specify **forward** if you want the frame destined to the specified interface to be forwarded. Specify **discard** if you want the frame to be discarded. • (Optional) For *interface-id*, specify the interface on which the address can be reached.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To disable the frame forwarding ability, use the **no bridge** *bridge-group* **address** *mac-address* global configuration command.

This example shows how to forward a frame with MAC address 0800.cb00.45e9 through an interface in bridge group 1:

```
Switch(config)# bridge 1 address 0800.cb00.45e9 forward gigabitethernet0/1
```

Adjusting Spanning-Tree Parameters

You might need to adjust certain spanning-tree parameters if the default values are not suitable for your switch configuration. Parameters affecting the entire spanning tree are configured with variations of the **bridge** global configuration command. Interface-specific parameters are configured with variations of the **bridge-group** interface configuration command.

You can adjust spanning-tree parameters by performing any of the tasks in these sections:

- Changing the Switch Priority, page 26-8
- Changing the Interface Priority, page 26-8
- Assigning a Path Cost, page 26-9
- Adjusting BPDU Intervals, page 26-10
- Disabling the Spanning Tree on an Interface, page 26-12

Note Only network administrators with a good understanding of how switches and STP function should make adjustments to spanning-tree parameters. Poorly planned adjustments can have a negative impact on performance. A good source on switching is the IEEE 802.1d specification; for more information, refer to the “References and Recommended Reading” appendix in the *Cisco IOS Configuration Fundamentals Command Reference*.

Changing the Switch Priority

You can globally configure the priority of an individual switch when two switches tie for position as the root switch, or you can configure the likelihood that a switch will be selected as the root switch. This priority is determined by default; however, you can change it.

Beginning in privileged EXEC mode, follow these steps to change the switch priority:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **priority** *number*	Change the priority of the switch. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *number*, enter a number from 0 to 65535. The default is 32768. The lower the number, the more likely the switch will be chosen as the root.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

No **no** form of this command exists. To return to the default setting, use the **bridge** *bridge-group* **priority** *number* global configuration command, and set the priority to the default value. To change the priority on an interface, use the **bridge-group priority** interface configuration command (described in the next section).

This example shows how to set the switch priority to 100 for bridge group 10:

```
Switch(config)# bridge 10 priority 100
```

Changing the Interface Priority

You can change the priority for an interface. When two switches tie for position as the root switch, you configure an interface priority to break the tie. The switch with the lowest interface value is elected.

Beginning in privileged EXEC mode, follow these steps to change the interface priority:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to set the priority.
Step 3	**bridge-group** *bridge-group* **priority** *number*	Change the priority of an interface. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *number*, enter a number from 0 to 255. The lower the number, the more likely that the interface on the switch will be chosen as the root. The default is 128.
Step 4	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 5	**show running-config**	Verify your entry.
Step 6	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

No **no** form of this command exists. To return to the default setting, use the **bridge-group** *bridge-group* **priority** *number* interface configuration command.

This example shows how to change the priority of an interface to 20 in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge-group 10 priority 20
```

Assigning a Path Cost

Each interface has a path cost associated with it. By convention, the path cost is 1000/data rate of the attached LAN, in Mbps.

Beginning in privileged EXEC mode, follow these steps to assign a path cost:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface to set the path cost.
Step 3	**bridge-group** *bridge-group* **path-cost** *cost*	Assign the path cost of an interface. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *cost*, enter a number from 1 to 65536. The higher the value, the higher the cost. – For 10 Mbps, the default path cost is 100. – For 100 Mbps, the default path cost is 19. – For 1000 Mbps, the default path cost is 4.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entry.
Step 6	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To return to the default path cost, use the **no bridge-group** *bridge-group* **path-cost** *cost* interface configuration command.

This example shows how to change the path cost on an interface to 10 in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge-group 10 path-cost 20
```

Adjusting BPDU Intervals

You can adjust BPDU intervals as described in these sections:

- Adjusting the Interval between Hello BPDUs
- Defining the Forward Delay Interval
- Defining the Maximum Idle Interval

Note Each switch in a spanning tree adopts the interval between hello BPDUs, the forward delay interval, and the maximum idle interval parameters of the root switch, regardless of what its individual configuration might be.

Adjusting the Interval between Hello BPDUs

Beginning in privileged EXEC mode, follow these step to adjust the interval between hello BPDUs:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **hello-time** *seconds*	Specify the interval between hello BPDUs. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 1 to 10. The default is 2 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge** *bridge-group* **hello-time** global configuration command.

This example shows how to change the hello interval to 5 seconds in bridge group 10:

```
Switch(config)# bridge 10 hello-time 5
```

Changing the Forward-Delay Interval

The forward-delay interval is the amount of time spent listening for topology change information after an interface has been activated for switching and before forwarding actually begins.

Beginning in privileged EXEC mode, follow these steps to change the forward-delay interval:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **forward-time** *seconds*	Specify the forward-delay interval. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 10 to 200. The default is 20 seconds.

	Command	Purpose
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge** *bridge-group* **forward-time** *seconds* global configuration command.

This example shows how to change the forward-delay interval to 10 seconds in bridge group 10:

```
Switch(config)# bridge 10 forward-time 10
```

Changing the Maximum-Idle Interval

If a switch does not hear BPDUs from the root switch within a specified interval, it recomputes the spanning-tree topology.

Beginning in privileged EXEC mode, follow these steps to change the maximum-idle interval (maximum aging time):

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**bridge** *bridge-group* **max-age** *seconds*	Specify the interval the switch waits to hear BPDUs from the root switch. • For *bridge-group*, specify the bridge group number. The range is 1 to 255. • For *seconds*, enter a number from 10 to 200. The default is 30 seconds.
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show running-config**	Verify your entry.
Step 5	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To return to the default setting, use the **no bridge** *bridge-group* **max-age** global configuration command.

This example shows how to change the maximum-idle interval to 30 seconds in bridge group 10:

```
Switch(config)# bridge 10 max-age 30
```

Disabling the Spanning Tree on an Interface

When a loop-free path exists between any two switched subnetworks, you can prevent BPDUs generated in one switching subnetwork from impacting devices in the other switching subnetwork, yet still permit switching throughout the network as a whole. For example, when switched LAN subnetworks are separated by a WAN, BPDUs can be prevented from traveling across the WAN link.

Beginning in privileged EXEC mode, follow these steps to disable spanning tree on an interface:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and specify the interface ID.
Step 3	**bridge-group** *bridge-group* **spanning-disabled**	Disable spanning tree on the interface. For *bridge-group*, specify the bridge group number. The range is 1 to 255.
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show running-config**	Verify your entry.
Step 6	**copy running-config startup-config**	(Optional) Save your entry in the configuration file.

To re-enable spanning tree on the interface, use the **no bridge-group** *bridge-group* **spanning-disabled** interface configuration command.

This example shows how to disable spanning tree on an interface in bridge group 10:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# bridge group 10 spanning-disabled
```

Monitoring and Maintaining the Network

To monitor and maintain the network, use one or more of the privileged EXEC commands in Table 26-2:

Table 26-2 Fallback Bridging Commands for Monitoring and Maintaining the Network

Command	Purpose
clear bridge *bridge-group*	Removes any learned entries from the forwarding database and clears the transmit and receive counts for any statically configured entries.
show bridge [*bridge-group*]	Displays details about the bridge group.
show bridge [*bridge-group*] [*interface-id*] [*address*] [**group**] [**verbose**]	Displays classes of entries in the bridge forwarding database.

CHAPTER 27

Troubleshooting

This chapter describes how to identify and resolve software problems related to the IOS software. Depending on the nature of the problem, you can use the command-line interface (CLI) or the Cluster Management Suite (CMS) to identify and solve problems.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the *Cisco IOS Command Summary for Release 12.1*.

This chapter consists of these sections:

- Using Recovery Procedures, page 27-1
- Preventing Autonegotiation Mismatches, page 27-10
- Diagnosing Connectivity Problems, page 27-11
- Using Debug Commands, page 27-14
- Using the show forward Command, page 27-15
- Using the crashinfo File, page 27-17

Using Recovery Procedures

These recovery procedures require that you have physical access to the switch:

- Recovering from Corrupted Software, page 27-2
- Recovering from a Lost or Forgotten Password, page 27-3
- Recovering from a Command Switch Failure, page 27-7
- Recovering from Lost Member Connectivity, page 27-10

Recovering from Corrupted Software

Switch software can be corrupted during an upgrade, by downloading the wrong file to the switch, and by deleting the image file. In all of these cases, the switch does not pass the power-on self-test (POST), and there is no connectivity.

This procedure uses the XMODEM Protocol to recover from a corrupt or wrong image file. There are many software packages that support the XMODEM protocol, and this procedure is largely dependent on the emulation software you are using.

Step 1 Connect a PC with terminal-emulation software supporting the XMODEM Protocol to the switch console port.

Step 2 Set the line speed on the emulation software to 9600 baud.

Step 3 Unplug the switch power cord.

Step 4 Press the **Mode** button, and at the same time, reconnect the power cord to the switch.

You can release the **Mode** button a second or two after the LED above port 1X goes off. Several lines of information about the software appear along with instructions:

```
The system has been interrupted prior to initializing the flash file system. The following
commands will initialize the flash file system, and finish loading the operating system
software:

flash_init
load_helper
boot
```

Step 5 Initialize the Flash file system:

```
switch: flash_init
```

Step 6 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

Step 7 Load any helper files:

```
switch: load_helper
```

Step 8 Start the file transfer by using the XMODEM protocol.

```
switch: copy xmodem: flash:image_filename.bin
```

Step 9 After the XMODEM request appears, use the appropriate command on the terminal-emulation software to start the transfer and to copy the software image into Flash memory.

Recovering from a Lost or Forgotten Password

The default configuration for Catalyst 3550 switches allows an end user with physical access to the switch to recover from a lost password by interrupting the boot process during power-on and by entering a new password.

Note On Catalyst 3550 Fast Ethernet switches, a system administrator can disable some of the functionality of this feature by allowing an end user to reset a password only by agreeing to return to the default configuration. If you are an end user trying to reset a password on a Catalyst 3550 Fast Ethernet switch and password recover has been disabled, a status message shows this during the recovery process.

Follow the steps in this procedure if you have forgotten or lost the switch password.

Step 1 Connect a terminal or PC with terminal-emulation software to the switch console port.

Step 2 Set the line speed on the emulation software to 9600 baud.

Step 3 Unplug the switch power cord.

Step 4 Press the **Mode** button, and at the same time, reconnect the power cord to the switch.

You can release the **Mode** button a second or two after the LED above port 1X turns off. Several lines of information about the software appear with instructions, informing you if the password recovery procedure has been disabled or not.

- If you see a message that begins with this:

  ```
  The system has been interrupted prior to initializing the flash file system. The
  following commands will initialize the flash file system
  ```

  proceed to the "Password Recovery with Password Recovery Enabled" section on page 27-3, and follow the steps.

- If you see a message that begins with this:

  ```
  The password-recovery mechanism has been triggered, but
  is currently disabled.
  ```

  proceed to the "Procedure with Password Recovery Disabled" section on page 27-5, and follow the steps.

Password Recovery with Password Recovery Enabled

If the password-recovery mechanism is enabled, this message appears:

```
The system has been interrupted prior to initializing the flash file system. The following
commands will initialize the flash file system, and finish loading the operating system
software:

flash_init
load_helper
boot
```

Step 1 Initialize the Flash file system:

```
switch: flash_init
```

Step 2 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

Step 3 Load any helper files:

```
switch: load_helper
```

Step 4 Display the contents of Flash memory:

```
switch: dir flash:
```

The switch file system is displayed:

```
Directory of flash:
   13  drwx         192   Mar 01 1993 22:30:48  c3550-i5q3l2-mz-121-0.0.53
   11  -rwx        5825   Mar 01 1993 22:31:59  config.text
   17  -rwx          27   Mar 01 1993 22:30:57  env_vars
    5  -rwx          90   Mar 01 1993 22:30:57  system_env_vars
   18  -rwx         720   Mar 01 1993 02:21:30  vlan.dat

16128000 bytes total (10003456 bytes free)
```

Step 5 Rename the configuration file to config.text.old.

This file contains the password definition.

```
switch: rename flash:config.text flash:config.text.old
```

Step 6 Boot the system:

```
switch: boot
```

You are prompted to start the setup program. Enter **N** at the prompt:

```
Continue with the configuration dialog? [yes/no]: N
```

Step 7 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
```

Step 8 Rename the configuration file to its original name:

```
Switch# rename flash:config.text.old flash:config.text
```

Step 9 Copy the configuration file into memory:

```
Switch# copy flash:config.text system:running-config
Source filename [config.text]?
Destination filename [running-config]?
```

Press **Return** in response to the confirmation prompts.

The configuration file is now reloaded, and you can change the password.

Step 10 Enter global configuration mode:

```
Switch# configure terminal
```

Step 11 Change the password:

```
Switch (config)# enable secret password
```

The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

Step 12 Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

Step 13 Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.

Note This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan** *vlan-id* global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

Procedure with Password Recovery Disabled

If the password-recovery mechanism is disabled, this message appears:

```
The password-recovery mechanism has been triggered, but
is currently disabled.  Access to the boot loader prompt
through the password-recovery mechanism is disallowed at
this point.  However, if you agree to let the system be
reset back to the default system configuration, access
to the boot loader prompt can still be allowed.

Would you like to reset the system back to the default configuration (y/n)?
```

Caution Returning the switch to the default configuration results in the loss of all existing configurations. We recommend that you contact your system administrator to verify if there are backup switch and VLAN configuration files.

- If you enter **n** (no), the normal boot process continues as if the **Mode** button had not been pressed; you cannot access the boot loader prompt, and you cannot enter a new password. You see the message:

  ```
  Press Enter to continue........
  ```

- If you enter **y** (yes), the configuration file in Flash memory and the VLAN database file are deleted. When the default configuration loads, you can reset the password.

Step 1 Elect to continue with password recovery and lose the existing configuration:

```
Would you like to reset the system back to the default configuration (y/n)? Y
```

Step 2 Load any helper files:

```
Switch: load_helper
```

Step 3 Display the contents of Flash memory:

```
switch: dir flash:
```

The switch file system is displayed:

```
Directory of flash:
   13  drwx         192   Mar 01 1993 22:30:48  c3550-i5q3l2-mz-121-0.0.53
   17  -rwx          27   Mar 01 1993 22:30:57  env_vars
    5  -rwx          90   Mar 01 1993 22:30:57  system_env_vars

16128000 bytes total (10003456 bytes free)
```

Step 4 Boot the system:

```
Switch: boot
```

You are prompted to start the setup program. To continue with password recovery, enter **N** at the prompt:

```
    Continue with the configuration dialog? [yes/no]: N
```

Step 5 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
```

Step 6 Enter global configuration mode:

```
Switch# configure terminal
```

Step 7 Change the password:

```
Switch (config)# enable secret password
```

The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

Step 8 Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

Step 9 Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.

Note This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan** *vlan-id* global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

Step 10 You must now reconfigure the switch. If the system administrator has the backup switch and VLAN configuration files available, you should use those.

Recovering from a Command Switch Failure

This section describes how to recover from a failed command switch. You can configure a redundant command switch group by using the Hot Standby Router Protocol (HSRP). For more information, see Chapter 5, "Clustering Switches" and Chapter 23, "Configuring HSRP."

Note HSRP is the preferred method for supplying redundancy to a cluster.

If you have not configured a standby command switch, and your command switch loses power or fails in some other way, management contact with the member switches is lost, and you must install a new command switch. However, connectivity between switches that are still connected is not affected, and the member switches forward packets as usual. You can manage the members as standalone switches through the console port or, if they have IP addresses, through the other management interfaces.

You can prepare for a command switch failure by assigning an IP address to a member switch or another switch that is command-capable, making a note of the command-switch password, and cabling your cluster to provide redundant connectivity between the member switches and the replacement command switch. This section describes two solutions for replacing a failed command switch:

- Replacing a failed command switch with a cluster member
- Replacing a failed command switch with another switch

For information on command-capable switches, refer to the release notes.

Replacing a Failed Command Switch with a Cluster Member

To replace a failed command switch with a command-capable member in the same cluster, follow these steps:

Step 1 Disconnect the command switch from the member switches, and physically remove it from the cluster.

Step 2 Insert the member switch in place of the failed command switch, and duplicate its connections to the cluster members.

Step 3 Start a CLI session on the new command switch.

You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, refer to the switch hardware installation guide.

Step 4 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
Switch#
```

Step 5 Enter the password of the *failed command switch*.

Step 6 Enter global configuration mode.

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
```

Step 7 Remove the member switch from the cluster.

```
Switch(config)# no cluster commander-address
```

Step 8 Return to privileged EXEC mode.

```
Switch(config)# end
Switch#
```

Step 9 Use the setup program to configure the switch IP information. This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.

```
Switch# setup
         --- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:
```

Step 10 Enter **Y** at the first prompt.

The prompts in the setup program vary depending on the member switch you selected to be the command switch:

```
Continue with configuration dialog? [yes/no]: y
```

or

```
Configuring global parameters:
```

If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.

Step 11 Respond to the questions in the setup program.

When prompted for the host name, recall that on a command switch, the host name is limited to 28 characters; on a member switch to 31 characters. Do not use *-n*, where *n* is a number, as the last characters in a host name for any switch.

When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

Step 12 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.

Step 13 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.

Step 14 When prompted, assign a name to the cluster, and press **Return**.

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

Step 15 After the initial configuration displays, verify that the addresses are correct.

Step 16 If the displayed information is correct, enter **Y**, and press **Return**.

If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.

Step 17 Start your browser, and enter the IP address of the new command switch.

Step 18 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.

Replacing a Failed Command Switch with Another Switch

To replace a failed command switch with a switch that is command-capable but not part of the cluster, follow these steps:

Step 1 Insert the new switch in place of the failed command switch, and duplicate its connections to the cluster members.

Step 2 Start a CLI session on the new command switch.

You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, refer to the switch hardware installation guide.

Step 3 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
Switch#
```

Step 4 Enter the password of the *failed command switch.*

Step 5 Use the setup program to configure the switch IP information.

This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.

```
Switch# setup
         --- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:
```

Step 6 Enter **Y** at the first prompt.

The prompts in the setup program vary depending on the switch you selected to be the command switch:

```
Continue with configuration dialog? [yes/no]: y
```

or

```
Configuring global parameters:
```

If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.

Step 7 Respond to the questions in the setup program.

When prompted for the host name, recall that on a command switch, the host name is limited to 28 characters. Do not use *-n*, where *n* is a number, as the last characters in a host name for any switch.

When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

Step 8 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.

Step 9 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.

Step 10 When prompted, assign a name to the cluster, and press **Return**.

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

Step 11 When the initial configuration displays, verify that the addresses are correct.

Step 12 If the displayed information is correct, enter **Y**, and press **Return**.

If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.

Step 13 Start your browser, and enter the IP address of the new command switch.

Step 14 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.

Recovering from Lost Member Connectivity

Some configurations can prevent the command switch from maintaining contact with member switches. If you are unable to maintain management contact with a member, and the member switch is forwarding packets normally, check for these conflicts:

- A member switch (Catalyst 3550, Catalyst 3500 XL, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) cannot connect to the command switch through a port that is defined as a network port.
- Catalyst 3500 XL, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 member switches must connect to the command switch through a port that belongs to the same management VLAN.
- A member switch (Catalyst 3550, Catalyst 2950, Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) connected to the command switch through a secured port can lose connectivity if the port is disabled because of a security violation.

Preventing Autonegotiation Mismatches

The IEEE 802.3AB autonegotiation protocol manages the switch settings for speed (10 Mbps, 100 Mbps, and 1000 Mbps excluding GBIC ports) and duplex (half or full). There are situations when this protocol can incorrectly align these settings, reducing performance. A mismatch occurs under these circumstances:

- A manually-set speed or duplex parameter is different from the manually set speed or duplex parameter on the connected port.
- A port is set to autonegotiate, and the connected port is set to full duplex with no autonegotiation.

To maximize switch performance and ensure a link, follow one of these guidelines when changing the settings for duplex and speed:

- Let both ports autonegotiate both speed and duplex.
- Manually set the speed and duplex parameters for the ports on both ends of the connection.

Note If a remote device does not autonegotiate, configure the duplex settings on the two ports to match. The speed parameter can adjust itself even if the connected port does not autonegotiate.

Diagnosing Connectivity Problems

This section describes how to troubleshoot connectivity problems:

- Understanding Ping, page 27-11
- Executing Ping, page 27-11
- Understanding IP Traceroute, page 27-12
- Executing IP Traceroute, page 27-13

Understanding Ping

The switch supports IP ping, which you can use to test connectivity to remote hosts. Ping sends an echo request packet to an address and waits for a reply. Ping returns one of these responses:

- Normal response—The normal response (*hostname* is alive) occurs in 1 to 10 seconds, depending on network traffic.
- Destination does not respond—If the host does not respond, a *no-answer* message is returned.
- Unknown host—If the host does not exist, an *unknown host* message is returned.
- Destination unreachable—If the default gateway cannot reach the specified network, a *destination-unreachable* message is returned.
- Network or host unreachable—If there is no entry in the route table for the host or network, a *network or host unreachable* message is returned.

Executing Ping

If you attempt to ping a host in a different IP subnetwork, you must define a static route to the network or have IP routing configured to route between those subnets. For more information, see Chapter 22, "Configuring IP Unicast Routing."

IP routing is disabled by default on all switches. If you need to enable or configure IP routing, see Chapter 22, "Configuring IP Unicast Routing."

Beginning in privileged EXEC mode, follow this step to ping another device on the network from the switch:

	Command	Purpose
Step 1	**ping ip** *host* | *address*	Ping a remote host through IP or by supplying the host name or network address.

Note Though other protocol keywords are available with the **ping** command, they are not supported in this release.

This example shows how to ping an IP host:

```
Switch# ping 172.20.52.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echoes to 172.20.52.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
Switch#
```

Table 27-1 describes the possible ping character output.

Table 27-1 Ping Output Display Characters

Character	Description
!	Each exclamation point means receipt of a reply.
.	Each period means the network server timed out while waiting for a reply.
U	A destination unreachable error PDU was received.
C	A congestion experienced packet was received.
I	User interrupted test.
?	Unknown packet type.
&	Packet lifetime exceeded.

To terminate a ping session, enter the escape sequence (**Ctrl-^ X** by default). You enter the default by simultaneously pressing and releasing the **Ctrl**, **Shift**, and **6** keys, and then pressing the **X** key.

Understanding IP Traceroute

You can use IP traceroute to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer (Layer 3) devices, such as routers, that the traffic passes through on the way to the destination.

Your switches can participate as the source or destination of the **traceroute** privileged EXEC command and might or might not appear as a hop in the **traceroute** command output. If the switch is the destination of the traceroute, it is displayed as the final destination in the traceroute output. Intermediate switches do not show up in the traceroute output if they are only bridging the packet from one port to another within the same VLAN. However, if the intermediate switch is a multilayer switch that is routing a particular packet, this switch shows up as a hop in the traceroute output.

The **traceroute** privileged EXEC command uses the Time To Live (TTL) field in the IP header to cause routers and servers to generate specific return messages. Traceroute starts by sending a User Datagram Protocol (UDP) datagram to the destination host with the TTL field set to 1. If a router finds a TTL value of 1 or 0, it drops the datagram and sends back an Internet Control Message Protocol (ICMP) time-to-live-exceeded message to the sender. Traceroute determines the address of the first hop by examining the source address field of the ICMP time-to-live-exceeded message.

To identify the next hop, traceroute sends a UDP packet with a TTL value of 2. The first router decrements the TTL field by 1 and sends the datagram to the next router. The second router sees a TTL value of 1, discards the datagram, and returns the time-to-live-exceeded message to the source. This process continues until the TTL is incremented to a value large enough for the datagram to reach the destination host (or until the maximum TTL is reached).

To determine when a datagram reaches its destination, traceroute sets the UDP destination port number in the datagram to a very large value that the destination host is unlikely to be using. When a host receives a datagram destined to itself containing a destination port number that is unused locally, it sends an ICMP port unreachable error to the source. Because all errors except port unreachable errors come from intermediate hops, the receipt of a port unreachable error means this message was sent by the destination.

Executing IP Traceroute

Beginning in privileged EXEC mode, follow this step to trace the path packets take through the network:

	Command	Purpose
Step 1	**traceroute ip** *host*	Trace the path packets take through the network by using IP.

Note Though other protocol keywords are available with the **traceroute** privileged EXEC command, they are not supported in this release.

This example shows how to perform a **traceroute** to an IP host:

```
Switch# traceroute ip 171.9.15.10

Type escape sequence to abort.
Tracing the route to 171.69.115.10

  1 172.2.52.1 0 msec 0 msec 4 msec
  2 172.2.1.203 12 msec 8 msec 0 msec
  3 171.9.16.6 4 msec 0 msec 0 msec
  4 171.9.4.5 0 msec 4 msec 0 msec
  5 171.9.121.34 0 msec 4 msec 4 msec
  6 171.9.15.9 120 msec 132 msec 128 msec
  7 171.9.15.10 132 msec 128 msec 128 msec
Switch#
```

The display shows the hop count, IP address of the router, and the round-trip time in milliseconds for each of the three probes that are sent.

Table 27-2 Traceroute Output Display Characters

Character	Description
*	The probe timed out.
?	Unknown packet type.
A	Administratively unreachable. Usually, this output means that an access list is blocking traffic.
H	Host unreachable.
N	Network unreachable.
P	Protocol unreachable.
Q	Source quench.
U	Port unreachable.

To terminate a trace in progress, enter the escape sequence (**Ctrl-^ X** by default). You enter the default by simultaneously pressing and releasing the **Ctrl**, **Shift**, and **6** keys, and then pressing the **X** key.

Using Debug Commands

This section explains how you use **debug** commands to diagnose and resolve internetworking problems. It contains this information:

- Enabling Debugging on a Specific Feature, page 27-14
- Enabling All-System Diagnostics, page 27-15
- Redirecting Debug and Error Message Output, page 27-15

Caution Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. It is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system use.

Note For complete syntax and usage information for the Catalyst 3550-specific **debug** commands, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

Enabling Debugging on a Specific Feature

All **debug** commands are entered in privileged EXEC mode, and most **debug** commands take no arguments. For example, beginning in privileged EXEC mode, enter this command to enable the debugging for Switch Port Analyzer:

```
Switch# debug span-session
```

The switch continues to generate output until you enter the **no** form of the command.

If you enable a **debug** command and no output is displayed, consider these possibilities:

- The switch might not be properly configured to generate the type of traffic you want to monitor. Use the **show running-config** command to check its configuration.
- Even if the switch is properly configured, it might not generate the type of traffic you want to monitor during the particular period that debugging is enabled. Depending on the feature you are debugging, you can use commands such as the TCP/IP **ping** command to generate network traffic.

To disable debugging of SPAN, enter this command in privileged EXEC mode:

```
Switch# no debug span-session
```

Alternately, in privileged EXEC mode, you can enter the **undebug** form of the command:

```
Switch# undebug span-session
```

To display the state of each debugging option, enter this command in privileged EXEC mode:

```
Switch# show debugging
```

Enabling All-System Diagnostics

Beginning in privileged EXEC mode, enter this command to enable all-system diagnostics:

```
Switch# debug all
```

Caution Because debugging output takes priority over other network traffic, and because the **debug all** privileged EXEC command generates more output than any other **debug** command, it can severely diminish switch performance or even render it unusable. In virtually all cases, it is best to use more specific **debug** commands.

The **no debug all** privileged EXEC command disables all diagnostic output. Using the **no debug all** command is a convenient way to ensure that you have not accidentally left any **debug** commands enabled.

Redirecting Debug and Error Message Output

By default, the network server sends the output from **debug** commands and system error messages to the console. If you use this default, you can use a virtual terminal connection to monitor debug output instead of connecting to the console port.

Possible destinations include the console, virtual terminals, internal buffer, and UNIX hosts running a syslog server. The syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX and its derivatives.

Note Be aware that the debugging destination you use affects system overhead. Logging messages to the console produces very high overhead, whereas logging messages to a virtual terminal produces less overhead. Logging messages to a syslog server produces even less, and logging to an internal buffer produces the least overhead of any method.

For more information about system message logging, see Chapter 17, “Configuring System Message Logging.”

Using the show forward Command

The output from the **show forward** privileged EXEC command provides some useful information about the disposition of a packet entering an interface. Depending upon the parameters entered about the packet, the output provides lookup table results, maps and masks used to calculate forwarding destinations, bitmaps, and egress information.

Note For more syntax and usage information for the **show forward** command, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release.

This is an example of the output from the **show forward** privileged EXEC command for Fast Ethernet port 8, where VLAN ID, source and destination MAC addresses, and source and destination IP addresses were provided.

```
Switch# show forward fa0/8 vlan 8 0000.1111.2222 0022.3355.9800 ip 8.8.8.10 4.4.4.33 255
signature:00000007, comparison ind:10, control info:2000941A control map:00000000
vlan:8, vlanid entry:000C0012 00000000 00000000 04620000
vlan:8, vlanid entry:000C0012 00000000 00000000 04620000

lookup  key                                 bk adata    rawoff secoff sec
qos     940808080A04040421 800000000000FF0000 0  00000000 006304 004064  4
acl     940808080A04040421 800000000000FF0000 1  00000082 045408 002016  1
learn   187008000011112222 801008002233559800 0  80010003 002176 002176  0
forw    187008000011112222 801008002233559800 1  40020000 043328 010560  5

bridgeDestMap:  00000000 00000000 0000FFFF FFFFFFC7
vlanMask:       00000000 00000000 0000FFFF FFFFFE7F
portMask:       00000000 00000000 00000000 00000080
sourceMask:     00000000 00000000 00000000 00000000
globalMap:      00000000 00000000 00000000 00000000
globalMask:     00000000 00000000 0002FFFF EFFFFC03
forwMap:        00000000 00000000 00000000 00000100

frame notifies:
src u_dat vlan fl q-map
2   00    8    00 00000000 00000000 00000000 00000100
Egress q 8
signature:00000007, comparison ind:10, control info:2000941A control map:00000000
vlan:8, vlanid entry:000C0012 00000000 00000000 04620000
FastEthernet0/9 vlan 8, dst 0022.3355.9800 src 0000.1111.2222, cos 0x0, dscp 0x0
```

Much of this information is useful mainly for Technical Support personnel, who have access to detailed information about the switch application-specific integrated circuits (ASICs). However, you can look at the *Egress q* section to get information about the output interface. There is an egress section for each separate destination port. The important information is in the line containing the name of the output interface, output VLAN ID, and rewritten destination MAC address for the frame. The example shows that the output interface is Fast Ethernet port 9 and the output VLAN is VLAN 8 and provides the rewritten source and destination MAC address for the frame.

If the output interface is a trunk port that needs to transmit multiple copies of the frame on different VLANs (for example, for IP multicast frames), several lines might contain the same output interface name, but different output VLANs. If output security access control lists (ACLs) are present, it is possible that one or more of these *Egress q* sections will not contain a line listing an output port. This happens when the output ACL denies the packet.

When the CPU is one of the destinations for a packet, a *Cpu q* section is displayed, followed by a queue name. This name should correspond to one of the queue names in the output from the **show controllers cpu-interface** privileged EXEC command, where statistics are displayed for the number of packets received at each queue.

This is an example of the *Cpu q* section display:

```
Cpu q:100 - routing queue
```

Using the crashinfo File

The crashinfo file saves information that helps Cisco technical support representatives to debug problems that caused the IOS image to fail (crash). The switch writes the crash information to the console at the time of the failure, and the file is created the next time you boot the IOS image after the failure (instead of while the system is failing).

The information in the file includes the IOS image name and version that failed, a dump of the processor registers, and a stack trace. You can provide this information to the Cisco technical support representative by using the **show tech-support** privileged EXEC command.

All crashinfo files are kept in this directory on the Flash file system:

flash:/crashinfo/crashinfo_*n* where *n* is a sequence number.

Each new crashinfo file that is created uses a sequence number that is larger than any previously-existing sequence number, so the file with the largest sequence number describes the most recent failure. Version numbers are used instead of a timestamp because the Catalyst 3550 switches do not include a real-time clock. You cannot change the name of the file that the system will use when it creates the file. However, after the file is created, you can use the **rename** privileged EXEC command to rename it, but the contents of the renamed file will not be displayed by the **show stacks** or the **show tech-support** privileged EXEC command.

You can display the most recent crashinfo file (that is, the file with the highest sequence number at the end of its filename) by entering the **show stacks** or the **show tech-support** privileged EXEC command. You also can access the file by using any command that can copy or display files, such as the **more** or the **copy** privileged EXEC command.

Supported MIBs

This appendix lists the supported management information base (MIBs) for this release. It contains these sections:

- MIB List, page A-1
- Using FTP to Access the MIB Files, page A-2

MIB List

- RFC 1213
- IF-MIB
- CISCO-CDP-MIB
- CISCO-IMAGE-MIB
- CISCO-FLASH-MIB
- OLD-CISCO-CHASSIS-MIB
- CISCO-PAGP-MIB
- CISCO-VTP-MIB
- CISCO-HSRP-MIB
- OLD-CISCO-TS-MIB
- BRIDGE-MIB (RFC1493)
- CISCO-VLAN-MEMBERSHIP-MIB
- CISCO-VLAN-IFINDEX-RELATIONSHIP-MIB
- CISCO-STACK-MIB (only a subset of the available MIB objects are implemented; not all objects are supported)
- RMON 1 MIB (only RMON etherStats, etherHistory, alarms, and events are supported)
- IGMP MIB
- PIM MIB
- CISCO-STP-EXTENSIONS-MIB
- OSPF-MIB (RFC 1253)
- IPMROUTE-MIB

- CISCO-MEMORY-POOL-MIB
- CISCO-RTTMON-MIB (subsystems supported: sub_rtt_rmon and sub_rtt_rmonlib)
- CISCO-PROCESS-MIB
- OLD-CISCO-SYS-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-MAC-NOTIFICATION-MIB
- CISCO-IGMP-FILTER-MIB

Using FTP to Access the MIB Files

You can obtain each MIB file by using this procedure:

Step 1 Use FTP to access the server **ftp.cisco.com**.

Step 2 Log in with the username **anonymous**.

Step 3 Enter your e-mail username when prompted for the password.

Step 4 At the `ftp>` prompt, change directories to **/pub/mibs/v1** and the **/pub/mibs/v2**.

Step 5 Use the **get** *MIB_filename* command to obtain a copy of the MIB file.

Note You can also access information about MIBs on the Cisco web site: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Working with the IOS File System, Configuration Files, and Software Images

This appendix describes how to manipulate the Flash file system, how to copy configuration files, and how to archive (upload and download) software images.

Note For complete syntax and usage information for the commands used in this chapter, refer to the *Catalyst 3550 Multilayer Switch Command Reference* for this release and the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This appendix consists of these sections:

- Working with the Flash File System, page B-1
- Working with Configuration Files, page B-8
- Working with Software Images, page B-19

Working with the Flash File System

The Flash file system on your switch provides several commands to help you manage software image and configuration files.

The Flash file system is a single Flash device on which you can store files. This Flash device is called *flash:*.

This section contains this information:

- Displaying Available File Systems, page B-2
- Setting the Default File System, page B-3
- Displaying Information about Files on a File System, page B-3
- Creating and Removing Directories, page B-4
- Copying Files, page B-4
- Deleting Files, page B-5
- Creating, Displaying, and Extracting tar Files, page B-6
- Displaying the Contents of a File, page B-8

Displaying Available File Systems

To display the available file systems on your switch, use the **show file systems** privileged EXEC command as shown in this example:

```
Switch# show file systems
File Systems:

     Size(b)     Free(b)      Type  Flags  Prefixes
*   16128000    11118592     flash     rw   flash:
    16128000    11118592   unknown     rw   zflash:
       32768       26363     nvram     rw   nvram:
           -           -   network     rw   tftp:
           -           -    opaque     rw   null:
           -           -    opaque     rw   system:
           -           -    opaque     ro   xmodem:
           -           -    opaque     ro   ymodem:
           -           -   network     rw   rcp:
           -           -   network     rw   ftp:
```

Table B-1 show file systems Field Descriptions

Field	Value
Size(b)	Amount of memory in the file system in bytes.
Free(b)	Amount of free memory in the file system in bytes.
Type	Type of file system. **flash**—The file system is for a Flash memory device. **nvram**—The file system is for a nonvolatile RAM (NVRAM) device. **opaque**—The file system is a locally generated *pseudo* file system (for example, the *system*) or a download interface, such as brimux. **unknown**—The file system is an unknown type.
Flags	Permission for file system. **ro**—read-only. **rw**—read/write. **wo**—write-only.
Prefixes	Alias for file system. **flash:**—Flash file system. **nvram:**—NVRAM. **null:**—Null destination for copies. You can copy a remote file to null to determine its size. **rcp:**—Remote Copy Protocol (RCP) network server. **system:**—Contains the system memory, including the running configuration. **tftp:**—Trivial File Transfer Protocol (TFTP) network server. **xmodem:**—Obtain the file from a network machine by using the XMODEM protocol. **ymodem:**—Obtain the file from a network machine by using the YMODEM protocol. **zflash:**—Read-only file decompression file system, which mirrors the contents of the Flash file system.

Setting the Default File System

You can specify the file system or directory that the system uses as the default file system by using the **cd** *filesystem:* privileged EXEC command. You can set the default file system to omit the *filesystem:* argument from related commands. For example, for all privileged EXEC commands that have the optional *filesystem:* argument, the system uses the file system specified by the **cd** command.

By default, the default file system is *flash:*.

You can display the current default file system as specified by the **cd** command by using the **pwd** privileged EXEC command.

Displaying Information about Files on a File System

You can view a list of the contents of a file system before manipulating its contents. For example, before copying a new configuration file to Flash memory, you might want to verify that the file system does not already contain a configuration file with the same name. Similarly, before copying a Flash configuration file to another location, you might want to verify its filename for use in another command.

To display information about files on a file system, use one of the privileged EXEC commands in Table B-2:

Table B-2 Commands for Displaying Information About Files

Command	Description
dir [/**all**] [*filesystem***:**][*filename*]	Display a list of files on a file system.
show file systems	Display more information about each of the files on a file system.
show file information *file-url*	Display information about a specific file.
show file descriptors	Display a list of open file descriptors. File descriptors are the internal representations of open files. You can use this command to see if another user has a file open.

Changing Directories and Displaying the Working Directory

Beginning in privileged EXEC mode, follow these steps to change directories and display the working directory.

	Command	Purpose
Step 1	**dir** *filesystem***:**	Display the directories on the specified file system. For *filesystem***:**, use **flash:** for the system board Flash device.
Step 2	**cd new_configs**	Change to the directory of interest. The command example shows how to change to the directory named *new_configs*.
Step 3	**pwd**	Display the working directory.

Creating and Removing Directories

Beginning in privileged EXEC mode, follow these steps to create and remove a directory:

	Command	Purpose
Step 1	**dir** *filesystem***:**	Display the directories on the specified file system. For *filesystem***:**, use **flash:** for the system board Flash device.
Step 2	**mkdir old_configs**	Create a new directory. The command example shows how to create the directory named *old_configs*. Directory names are case sensitive. Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons.
Step 3	**dir** *filesystem***:**	Verify your entry.

To delete a directory with all its files and subdirectories, use the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command.

Use the **/recursive** keyword to delete the named directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the name of the directory to be deleted. All the files in the directory and the directory are removed.

Caution When files and directories are deleted, their contents cannot be recovered.

Copying Files

To copy a file from a source to a destination, use the **copy** [**/erase**] *source-url destination-url* privileged EXEC command. For the source and destination URLs, you can use **running-config** and **startup-config** keyword shortcuts. For example, the **copy running-config startup-config** command saves the currently running configuration file to the NVRAM section of Flash memory to be used as the configuration during system initialization.

You can also copy to and from special file systems (**xmodem:**, **ymodem:**) as the source or destination for the file from a network machine that uses the XMODEM or YMODEM protocol.

Network file system URLs include **ftp:**, **rcp:**, and **tftp:** and have these syntaxes:

File Transfer Protocol (FTP)—**ftp:**[[*//username* [**:***password*]**@***location*]*/directory*]*/filename*

Remote Copy Protocol (RCP)—**rcp:**[[*//username***@***location*]*/directory*]*/filename*

Trivial File Transfer Protocol (TFTP)—**tftp:**[[*//location*]*/directory*]*/filename*

Local writable file systems include flash:.

Some invalid combinations of source and destination exist. Specifically, you cannot copy these combinations:

- From a running configuration to a running configuration
- From a startup configuration to a startup configuration
- From a device to the same device (for example, the **copy flash: flash:** command is invalid)

For specific examples of using the **copy** command with configuration files, see the "Working with Configuration Files" section on page B-8.

To copy software images either by downloading a new version or uploading the existing one, use the **archive download-sw** or the **archive upload-sw** privileged EXEC command. For more information, see the "Working with Software Images" section on page B-19.

Deleting Files

When you no longer need a file on a Flash memory device, you can permanently delete it. To delete a file or directory from a specified Flash device, use the **delete** [**/force**] [**/recursive**] [*filesystem***:**]*file-url* privileged EXEC command.

Use the **/recursive** keyword for deleting a directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

If you omit the *filesystem***:** option, the switch uses the default device specified by the **cd** command. For *file-url*, you specify the path (directory) and the name of the file to be deleted.

If you attempt to delete the file specified by the CONFIG_FILE or BOOT environment variable, the system prompts you to confirm the deletion. If you attempt to delete the last valid system image specified in the BOOT environment variable, the system prompts you to confirm the deletion.

Caution When files are deleted, their contents cannot be recovered.

This example shows how to delete the file *myconfig* from the default Flash memory device:

```
Switch# delete myconfig
```

Creating, Displaying, and Extracting tar Files

You can create a tar file and write files into it, list the files in a tar file, and extract the files from a tar file as described in the next sections.

Creating a tar File

To create a tar file and write files into it, use this privileged EXEC command:

archive tar /create *destination-url* **flash:**/*file-url*

For *destination-url*, specify the destination URL alias for the local or network file system and the name of the tar file to create. These options are supported:

- For the local Flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:**[[//*username*[**:***password*]**@***location*]/*directory*]/*tar-filename***.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rcp:**[[//*username***@***location*]/*directory*]/*tar-filename***.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:**[[//*location*]/*directory*]/*tar-filename***.tar**

The *tar-filename***.tar** is the tar file to be created.

For **flash:**/*file-url*, specify the location on the local Flash file system from which the new tar file is created. You can also specify an optional list of files or directories within the source directory to write to the new tar file. If none are specified, all files and directories at this level are written to the newly created tar file.

This example shows how to create a tar file. This command writes the contents of the *new-configs* directory on the local Flash device to a file named *saved.tar* on the TFTP server at 172.20.10.30:

```
Switch# archive tar /create tftp:172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a tar File

To display the contents of a tar file on the screen, use this privileged EXEC command:

archive tar /table *source-url*

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local Flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:**[[//*username*[**:***password*]**@***location*]/*directory*]/*tar-filename***.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rcp:**[[//*username***@***location*]/*directory*]/*tar-filename***.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:**[[//*location*]/*directory*]/*tar-filename***.tar**

The *tar-filename***.tar** is the tar file to display.

You can also limit the display of the files by specifying an optional list of files or directories after the tar file; then only these files are displayed. If none are specified, all files and directories are displayed.

This example shows how to display the contents of the *c3550-i5q3l2-mz.121-6.EA1.tar* file that is in Flash memory:

```
Switch# archive tar /table flash:c3550-i5q3l2-mz.121-6.EA1.tar
info (219 bytes)
c3550-i5q3l2-mz.121-6.EA1/ (directory)
c3550-i5q3l2-mz.121-6.EA1/html/ (directory)
c3550-i5q3l2-mz.121-6.EA1/html/foo.html (0 bytes)
c3550-i5q3l2-mz.121-6.EA1/c3550-i5q3l2-mz.121-6.EA1.bin (610856 bytes)
c3550-i5q3l2-mz.121-6.EA1/info (219 bytes)
info.ver (219 bytes)
```

This example shows how to display only the *c3550-i5q3l2-mz.121-6.EA1/html* directory and its contents:

```
Switch# archive tar /table flash:c3550-tv0-m.tar c3550-i5q3l2-mz.121-6.EA1/html
c3550-i5q3l2-mz.121-6.EA1/html/ (directory)
c3550-i5q3l2-mz.121-6.EA1/html/foo.html (0 bytes)
```

Extracting a tar File

To extract a tar file into a directory on the Flash file system, use this privileged EXEC command:

archive tar /xtract *source-url* **flash:**/*file-url*

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local Flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:**[[//*username*[**:***password*]**@***location*]/*directory*]/*tar-filename***.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rcp:**[[//*username***@***location*]/*directory*]/*tar-filename***.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:**[[//*location*]/*directory*]/*tar-filename***.tar**

The *tar-filename***.tar** is the tar file from which to extract files.

For **flash:**/*file-url*, specify the location on the local Flash file system into which the tar file is extracted. You can also specify an optional list of files or directories within the tar file for extraction. If none are specified, all files and directories are extracted.

This example shows how to extract the contents of a tar file located on the TFTP server at 172.20.10.30. This command extracts just the *new-configs* directory into the root directory on the local Flash file system. The remaining files in the *saved.tar* file are ignored.

```
Switch# archive tar /xtract tftp:/172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a File

To display the contents of any readable file, including a file on a remote file system, use the **more** [**/ascii** | **/binary** | **/ebcdic**] *file-url* privileged EXEC command:

This example shows how to display the contents of a configuration file on a TFTP server:

```
Switch# more tftp://serverA/hampton/savedconfig
!
! Saved configuration on server
!
version 11.3
service timestamps log datetime localtime
service linenumber
service udp-small-servers
service pt-vty-logging
!

<output truncated>
```

Working with Configuration Files

This section describes how to create, load, and maintain configuration files. Configuration files contain commands entered to customize the function of the Cisco IOS software. To better benefit from these instructions, your switch must contain a minimal configuration for interacting with the system software. You can create a basic configuration file by using the **setup** program or by entering the **setup** privileged EXEC command. For more information, see Chapter 4, “Assigning the Switch IP Address and Default Gateway.”

You can copy (*download*) configuration files from a TFTP, FTP, or RCP server to the running configuration or startup configuration of the switch. You might want to perform this for one of these reasons:

- To restore a backed-up configuration file.
- To use the configuration file for another switch. For example, you might add another switch to your network and want it to have a configuration similar to the original switch. By copying the file to the new switch, you can change the relevant parts rather than recreating the whole file.
- To load the same configuration commands on all the switches in your network so that all the switches have similar configurations.

You can copy (*upload*) configuration files from the switch to a file server by using TFTP, FTP, or RCP. You might perform this task to back up a current configuration file to a server before changing its contents so that you can later restore the original configuration file from the server.

The protocol you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- Guidelines for Creating and Using Configuration Files, page B-9
- Configuration File Types and Location, page B-9
- Creating a Configuration File By Using a Text Editor, page B-10
- Copying Configuration Files By Using TFTP, page B-10

- Copying Configuration Files By Using FTP, page B-12
- Copying Configuration Files By Using RCP, page B-16
- Clearing Configuration Information, page B-19

Guidelines for Creating and Using Configuration Files

Creating configuration files can aid in your switch configuration. Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration.

Use these guidelines when creating a configuration file:

- We recommend that you connect through the console port when using configuration files to configure the switch. If you configure the switch from a Telnet session, IP addresses are not changed, and ports and modules are not disabled.
- If no passwords have been set on the switch, you must set them on each switch by entering the **enable secret** *secret-password* global configuration command. Enter a blank line for this command. The password is saved in the configuration file as clear text.
- If passwords already exist, you cannot enter the **enable secret** *secret-password* global configuration command in the file because the password verification will fail. If you enter a password in the configuration file, the switch mistakenly attempts to execute the passwords as commands as it executes the file.

Note The **copy** {**ftp:** | **rcp:** | **tftp:**} **system:running-config** privileged EXEC command loads the configuration files on the switch as if you were entering the commands at the command line. The switch does not erase the existing running configuration before adding the commands. If a command in the copied configuration file replaces a command in the existing configuration file, the existing command is erased. For example, if the copied configuration file contains a different IP address in a particular command than the existing configuration, the IP address in the copied configuration is used. However, some commands in the existing configuration might not be replaced or negated. In this case, the resulting configuration file is a mixture of the existing configuration file and the copied configuration file, with the copied configuration file having precedence.

To restore a configuration file to an exact copy of a file stored on a server, copy the configuration file directly to the startup configuration (by using the **copy** {**ftp:** | **rcp:** | **tftp:**} **nvram:startup-config** privileged EXEC command), and reload the switch.

Configuration File Types and Location

Startup configuration files are used during system startup to configure the software. Running configuration files contain the current configuration of the software. The two configuration files can be different. For example, you might want to change the configuration for a short time period rather than permanently. In this case, you would change the running configuration but not save the configuration by using the **copy running-config startup-config** privileged EXEC command.

The running configuration is saved in DRAM; the startup configuration is stored in the NVRAM section of Flash memory.

Creating a Configuration File By Using a Text Editor

When creating a configuration file, you must list commands logically so that the system can respond appropriately. This is one method of creating a configuration file:

Step 1 Copy an existing configuration from a switch to a server.

For more information, see the "Downloading the Configuration File By Using TFTP" section on page B-11, the "Downloading a Configuration File By Using FTP" section on page B-13, or the "Downloading a Configuration File By Using RCP" section on page B-17.

Step 2 Open the configuration file in a text editor, such as vi or emacs on UNIX or Notepad on a PC.

Step 3 Extract the portion of the configuration file with the desired commands, and save it in a new file.

Step 4 Copy the configuration file to the appropriate server location. For example, copy the file to the TFTP directory on the workstation (usually /tftpboot on a UNIX workstation).

Step 5 Make sure the permissions on the file are set to world-read.

Copying Configuration Files By Using TFTP

You can configure the switch by using configuration files you create, download from another switch, or download from a TFTP server. You can copy (upload) configuration files to a TFTP server for storage.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using TFTP, page B-10
- Downloading the Configuration File By Using TFTP, page B-11
- Uploading the Configuration File By Using TFTP, page B-12

Preparing to Download or Upload a Configuration File By Using TFTP

Before you begin downloading or uploading a configuration file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the /etc/inetd.conf file contains this line:

  ```
  tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
  ```

  Make sure that the /etc/services file contains this line:

  ```
  tftp 69/udp
  ```

  Note You must restart the inetd daemon after modifying the /etc/inetd.conf and /etc/services files. To restart the daemon, either stop the inetd process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.

- Ensure that the configuration file to be downloaded is in the correct directory on the TFTP server (usually /tftpboot on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the configuration file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch** *filename* command, where *filename* is the name of the file you will use when uploading it to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading the Configuration File By Using TFTP

To configure the switch by using a configuration file downloaded from a TFTP server, follow these steps:

Step 1 Copy the configuration file to the appropriate TFTP directory on the workstation.

Step 2 Verify that the TFTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using TFTP" section on page B-10.

Step 3 Log into the switch through the console port or a Telnet session.

Step 4 Download the configuration file from the TFTP server to configure the switch.

Specify the IP address or host name of the TFTP server and the name of the file to download.

Use one of these privileged EXEC commands:

- **copy tftp:**[[[//*location*]/*directory*]/*filename*] **system:running-config**
- **copy tftp:**[[[//*location*]/*directory*]/*filename*] **nvram:startup-config**

The configuration file downloads, and the commands are executed as the file is parsed line-by-line.

This example shows how to configure the software from the file *tokyo-confg* at IP address 172.16.2.155:

```
Switch# copy tftp://172.16.2.155/tokyo-confg system:running-config
Configure using tokyo-confg from 172.16.2.155? [confirm] y
Booting tokyo-confg from 172.16.2.155:!!! [OK - 874/16000 bytes]
```

Uploading the Configuration File By Using TFTP

To upload a configuration file from a switch to a TFTP server for storage, follow these steps:

Step 1 Verify that the TFTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using TFTP" section on page B-10.

Step 2 Log into the switch through the console port or a Telnet session.

Step 3 Upload the switch configuration to the TFTP server. Specify the IP address or host name of the TFTP server and the destination filename.

Use one of these privileged EXEC commands:

- **copy system:running-config tftp:**[[[*//location*]*/directory*]*/filename*]
- **copy nvram:startup-config tftp:**[[[*//location*]*/directory*]*/filename*]

The file is uploaded to the TFTP server.

This example shows how to upload a configuration file from a switch to a TFTP server:

```
Switch# copy system:running-config tftp://172.16.2.155/tokyo-confg
Write file tokyo-confg on host 172.16.2.155? [confirm] y
#
Writing tokyo-confg!!! [OK]
```

Copying Configuration Files By Using FTP

You can copy configuration files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy a configuration file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **copy** command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept your FTP write request.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **copy** command if you want to specify only a username for that copy operation.

If the server has a directory structure, the configuration file is written to or copied from the directory associated with the username on the server. For example, if the configuration file resides in the home directory of a user on the server, specify that user's name as the remote username.

For more information, refer to the documentation for your FTP server.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using FTP, page B-13
- Downloading a Configuration File By Using FTP, page B-13
- Uploading a Configuration File By Using FTP, page B-15

Preparing to Download or Upload a Configuration File By Using FTP

Before you begin downloading or uploading a configuration file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username** *username* global configuration command during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a configuration file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using FTP:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode on the switch. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Change the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Change the default password.

	Command	Purpose
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**copy ftp:**[[[//[*username*[**:***password*]**@**]*location*]/*directory*] /*filename*] **system:running-config** or **copy ftp:**[[[//[*username*[**:***password*]**@**]*location*]/*directory*] /*filename*] **nvram:startup-config**	Using FTP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and to load and run those commands on the switch:

```
Switch# copy ftp://netadmin1:mypass@172.16.101.101/host1-confg system:running-config
Configure using host1-confg from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-confg:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by ftp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. The software copies the configuration file *host2-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the switch startup configuration.

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin1
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy ftp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-confg]? host2-confg
Configure using host2-confg from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-confg:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by ftp from
172.16.101.101
```

Uploading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using FTP:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Change the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Change the default password.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**copy system:running-config ftp:**[[[//[*username*[**:***password*]**@**]*location*]/*directory*] /*filename*] or **copy nvram:startup-config ftp:**[[[//[*username*[**:***password*]**@**]*location*]/*directory*] /*filename*]	Using FTP, store the switch running or startup configuration file to the specified location.

This example shows how to copy the running configuration file named *switch2-confg* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config ftp://netadmin1:mypass@172.16.101.101/switch2-confg
Write file switch2-confg on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server by using FTP to copy the file:

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin2
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy nvram:startup-config ftp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-confg]?
Write file switch2-confg on host 172.16.101.101?[confirm]
![OK]
```

Copying Configuration Files By Using RCP

The Remote Copy Protocol (RCP) provides another method of downloading, uploading, and copying configuration files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

The RCP requires a client to send a remote username with each RCP request to a server. When you copy a configuration file from the switch to a server, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip rcmd remote-username** *username* global configuration command if the command is configured.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For a successful RCP copy request, you must define an account on the network server for the remote username. If the server has a directory structure, the configuration file is written to or copied from the directory associated with the remote username on the server. For example, if the configuration file is in the home directory of a user on the server, specify that user's name as the remote username.

This section includes this information:

- Preparing to Download or Upload a Configuration File By Using RCP, page B-16
- Downloading a Configuration File By Using RCP, page B-17
- Uploading a Configuration File By Using RCP, page B-18

Preparing to Download or Upload a Configuration File By Using RCP

Before you begin downloading or uploading a configuration file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username** *username* global configuration command to be used during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the RCP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.

- When you upload a file to the RCP server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the .rhosts file for the remote user on the RCP server. For example, suppose that the switch contains these configuration lines:

  ```
  hostname Switch1
  ip rcmd remote-username User0
  ```

  If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

  ```
  Switch1.company.com Switch1
  ```

For more information, refer to the documentation for your RCP server.

Downloading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using RCP:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page B-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specify the remote username.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**copy rcp:**[[[//[*username*@]*location*]/*directory*]/*filename*] **system:running-config** or **copy rcp:**[[[//[*username*@]*location*]/*directory*]/*filename*] **nvram:startup-config**	Using RCP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and load and run those commands on the switch:

```
Switch# copy rcp://netadmin1@172.16.101.101/host1-confg system:running-config
Configure using host1-confg from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-confg:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by rcp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. Then it copies the configuration file *host2-confg* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the startup configuration:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin1
Switch(config)# end
Switch# copy rcp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-confg]? host2-confg
Configure using host2-confg from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-confg:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by rcp from
172.16.101.101
```

Uploading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using RCP:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using RCP" section on page B-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specify the remote username.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**copy system:running-config rcp:**[[[//[*username***@**]*location*]/*directory*]/*filename*] or **copy nvram:startup-config rcp:**[[[//[*username***@**]*location*]/*directory*]/*filename*]	Using RCP, copy the configuration file from a switch running or startup configuration file to a network server.

This example shows how to copy the running configuration file named *switch2-confg* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config rcp://netadmin1@172.16.101.101/switch2-confg
Write file switch-confg on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin2
Switch(config)# end
Switch# copy nvram:startup-config rcp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-confg]?
Write file switch2-confg on host 172.16.101.101?[confirm]
![OK]
```

Clearing Configuration Information

You can clear the configuration information from the startup configuration. If you reboot the switch with no startup configuration, the switch enters the setup program so that you can reconfigure the switch with all new settings.

Clearing the Startup Configuration File

To clear the contents of your startup configuration, use the **erase nvram:** or the **erase startup-config** privileged EXEC command.

Caution You cannot restore the startup configuration file after it has been deleted.

Deleting a Stored Configuration File

To delete a saved configuration from Flash memory, use the **delete flash:***filename* privileged EXEC command. Depending on the setting of the **file prompt** global configuration command, you might be prompted for confirmation before you delete a file. By default, the switch prompts for confirmation on destructive file operations. For more information about the **file prompt** command, refer to the *Cisco IOS Command Reference for Release 12.1*.

Caution You cannot restore a file after it has been deleted.

Working with Software Images

This section describes how to archive (download and upload) software image files, which contain the system software, IOS code, and the web management HTML files.

You download a switch image file from a TFTP, FTP, or RCP server to upgrade the switch software. You can replace the current image with the new one or keep the current image in Flash memory after a download.

You upload a switch image file to a TFTP, FTP, or RCP server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

The protocol you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- Image Location on the Switch, page B-20
- tar File Format of Images on a Server or Cisco.com, page B-20
- Copying Image Files By Using TFTP, page B-21
- Copying Image Files By Using FTP, page B-25
- Copying Image Files By Using RCP, page B-29

Note For a list of software images and the supported upgrade paths, refer to the release notes that shipped with your switch.

Image Location on the Switch

The IOS image is stored as a *.bin* file in a directory that shows the version number. A subdirectory contains the HTML files needed for web management. The image is stored on the system board Flash memory (flash:).

You can use the **show version** privileged EXEC command to see the software version that is currently running on your switch. In the display, check the line that begins with `System image file is...`. It shows the directory name in Flash memory where the image is stored.

You can also use the **dir** *filesystem***:** privileged EXEC command to see the directory names of other software images you might have stored in Flash memory.

tar File Format of Images on a Server or Cisco.com

Software images located on a server or downloaded from Cisco.com are provided in a tar file format, which contains these files:

- *info* file

  The info file is always at the beginning of the tar file and contains information about the files within it.
- IOS image
- Web management files needed by the HTTP server on the switch
- *info.ver* file

  The info.ver file is always at the end of the tar file and contains the same information as the info file. Because it is the last file in the tar file, its existence means that all files in the image have been downloaded.

This example shows the information contained in the info and info.ver files:

```
version_suffix: i5q3l2-121-6.EA1
version_directory: c3550-i5q3l2-mz.121-6.EA1
image_name: c3550-i5q3l2-mz.121-6.EA1.bin
ios_image_file_size: 3049472
total_image_file_size: 4551168
image_feature: LAYER_3|MIN_DRAM_MEG=64
image_family: C3550
info_end:
```

Table B-3 info and info.ver File Description

Field	Description
version_suffix	Specifies the IOS image version string suffix
version_directory	Specifies the directory where the IOS image and the HTML subdirectory are installed
image_name	Specifies the name of the IOS image within the tar file
ios_image_file_size	Specifies the IOS image size in the tar file, which is an approximate measure of how much Flash space is required to hold just the IOS image
total_image_file_size	Specifies the size of all the images (the IOS image and the HTML files) in the tar file, which is an approximate measure of how much Flash space is required to hold them
image_feature	Describes the core functionality of the image
image_family	Describes the family of products on which the software can be installed
image_min_dram	Specifies the minimum amount of DRAM needed to run this image

Copying Image Files By Using TFTP

You can download a switch image from a TFTP server or upload the image from the switch to a TFTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes; this uploaded image can be used for future downloads to the same or another switch of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using TFTP, page B-22
- Downloading an Image File By Using TFTP, page B-22
- Uploading an Image File By Using TFTP, page B-24

Preparing to Download or Upload an Image File By Using TFTP

Before you begin downloading or uploading an image file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the /etc/inetd.conf file contains this line:

  ```
  tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
  ```

  Make sure that the /etc/services file contains this line:

  ```
  tftp 69/udp
  ```

  Note You must restart the inetd daemon after modifying the /etc/inetd.conf and /etc/services files. To restart the daemon, either stop the inetd process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the image to be downloaded is in the correct directory on the TFTP server (usually /tftpboot on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the image file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch** *filename* command, where *filename* is the name of the file you will use when uploading the image to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading an Image File By Using TFTP

You can download a new image file and replace the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 3 to download a new image from a TFTP server and overwrite the existing image. To keep the current image, skip Step 3.

	Command	Purpose
Step 1		Copy the image to the appropriate TFTP directory on the workstation. Make sure the TFTP server is properly configured; see the "Preparing to Download or Upload an Image File By Using TFTP" section on page B-22.
Step 2		Log into the switch through the console port or a Telnet session.

	Command	Purpose
Step 3	**archive download-sw /overwrite /reload tftp:**[[*//location*]*/directory*]*/image-name***.tar**	Download the image file from the TFTP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in Flash with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For *//location*, specify the IP address of the TFTP server. • For */directory/image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 4	**archive download-sw /leave-old-sw /reload tftp:**[[*//location*]*/directory*]*/image-name***.tar**	Download the image file from the TFTP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For *//location*, specify the IP address of the TFTP server. • For */directory/image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the current running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image on the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using TFTP

You can upload an image from the switch to a TFTP server. You can later download this image to the switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to a TFTP server:

	Command	Purpose
Step 1		Make sure the TFTP server is properly configured; see the "Preparing to Download or Upload an Image File By Using TFTP" section on page B-22.
Step 1		Log into the switch through the console port or a Telnet session.
Step 2	**archive upload-sw tftp:**[[*//location*]*/directory*]*/image-name***.tar**	Upload the currently running switch image to the TFTP server. • For *//location*, specify the IP address of the TFTP server. • For */directory/image-name***.tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The *image-name***.tar** is the name of the software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using FTP

You can download a switch image from an FTP server or upload the image from the switch to an FTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the switch or another switch of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using FTP, page B-25
- Downloading an Image File By Using FTP, page B-26
- Uploading an Image File By Using FTP, page B-28

Preparing to Download or Upload an Image File By Using FTP

You can copy images files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy an image file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept the FTP write request from you.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.

If the server has a directory structure, the image file is written to or copied from the directory associated with the username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username** *username* global configuration command. This new name will be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username for that operation only.
- When you upload an image file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading an Image File By Using FTP

You can download a new image file and overwrite the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 7 to download a new image from an FTP server and overwrite the existing image. To keep the current image, skip Step 7.

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload an Image File By Using FTP" section on page B-25.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Change the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Change the default password.
Step 6	**end**	Return to privileged EXEC mode.

	Command	Purpose
Step 7	**archive download-sw /overwrite /reload ftp:**[[*//username*[**:***password*]**@***location*]*/directory*] */image-name***.tar**	Download the image file from the FTP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in Flash with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For *//username*[**:***password*], specify the username and password; these must be associated with an account on the FTP server. For more information, see the "Preparing to Download or Upload an Image File By Using FTP" section on page B-25. • For **@***location*, specify the IP address of the FTP server. • For *directory*/*image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 8	**archive download-sw /leave-old-sw /reload ftp:**[[*//username*[**:***password*]**@***location*]*/directory*] */image-name***.tar**	Download the image file from the FTP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For *//username*[**:***password*], specify the username and password. These must be associated with an account on the FTP server. For more information, see the "Preparing to Download or Upload an Image File By Using FTP" section on page B-25. • For **@***location*, specify the IP address of the FTP server. • For *directory*/*image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device, whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:/***file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using FTP

You can upload an image from the switch to an FTP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an FTP server:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the "Preparing to Download or Upload a Configuration File By Using FTP" section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	**ip ftp username** *username*	(Optional) Change the default remote username.
Step 5	**ip ftp password** *password*	(Optional) Change the default password.
Step 6	**end**	Return to privileged EXEC mode.
Step 7	**archive upload-sw ftp:**[[//[*username*[**:***password*]**@**]*location*]/*directory*]/*image-name***.tar**	Upload the currently running switch image to the FTP server. • For //*username***:***password*, specify the username and password. These must be associated with an account on the FTP server. For more information, see the "Preparing to Download or Upload an Image File By Using FTP" section on page B-25. • For **@***location*, specify the IP address of the FTP server. • For /*directory*/*image-name***.tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The *image-name***.tar** is the name of the software image to be stored on the server.

The **archive upload-sw** command builds an image file on the server by uploading these files in order: info, the IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using RCP

You can download a switch image from an RCP server or upload the image from the switch to an RCP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

This section includes this information:

- Preparing to Download or Upload an Image File By Using RCP, page B-29
- Downloading an Image File By Using RCP, page B-30
- Uploading an Image File By Using RCP, page B-32

Preparing to Download or Upload an Image File By Using RCP

RCP provides another method of downloading and uploading image files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

RCP requires a client to send a remote username on each RCP request to a server. When you copy an image from the switch to a server by using RCP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip rcmd remote-username** *username* global configuration command if the command is entered.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. If the server has a directory structure, the image file is written to or copied from the directory associated with the remote username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username** *username* global configuration command to be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and there is no need to set the RCP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.
- When you upload an image to the RCP to the server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the .rhosts file for the remote user on the RCP server. For example, suppose the switch contains these configuration lines:

  ```
  hostname Switch1
  ip rcmd remote-username User0
  ```

  If the switch IP address translates to *Switch1.company.com*, the .rhosts file for User0 on the RCP server should contain this line:

  ```
  Switch1.company.com Switch1
  ```

For more information, refer to the documentation for your RCP server.

Downloading an Image File By Using RCP

You can download a new image file and replace or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 6 to download a new image from an RCP server and overwrite the existing image. To keep the current image, skip Step 6.

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload an Image File By Using RCP" section on page B-29.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specify the remote username.

	Command	Purpose
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**archive download-sw /overwrite /reload rcp:**[[[//[*username***@**]*location*]/*directory*]/*image-name***.tar**]	Download the image file from the RCP server to the switch, and overwrite the current image. • The **/overwrite** option overwrites the software image in Flash with the downloaded image. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //*username*, specify the username. For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. For more information, see the "Preparing to Download or Upload an Image File By Using RCP" section on page B-29. • For **@***location*, specify the IP address of the RCP server. • For /*directory*/*image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 7	**archive download-sw /leave-old-sw /reload rcp:**[[[//[*username***@**]*location*]/*directory*]/*image-name***.tar**]	Download the image file from the RCP server to the switch, and keep the current image. • The **/leave-old-sw** option keeps the old software version after a download. • The **/reload** option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //*username*, specify the username. For the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the "Preparing to Download or Upload an Image File By Using RCP" section on page B-29. • For **@***location*, specify the IP address of the RCP server. • For /*directory*]/*image-name***.tar**, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

Note If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough room to install the new image an keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old software during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive** *filesystem***:**/*file-url* privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using RCP

You can upload an image from the switch to an RCP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature is available only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an RCP server:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the "Preparing to Download or Upload an Image File By Using RCP" section on page B-29.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	**configure terminal**	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	**ip rcmd remote-username** *username*	(Optional) Specify the remote username.

	Command	Purpose
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**archive upload-sw rcp:**[[[//[*username***@**]*location*]/*directory*]/*image-name***.tar**]	Upload the currently running switch image to the RCP server. • For //*username*, specify the username; for the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the "Preparing to Download or Upload an Image File By Using RCP" section on page B-29. • For **@***location*, specify the IP address of the RCP server. • For /*directory*]/*image-name***.tar**, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. • The *image-name***.tar** is the name of software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the IOS image, the HTML files, and info.ver. After these files are uploaded, the upload algorithm creates the tar file format.

Caution For the download and upload algorithms to operate properly, do *not* rename image names.

Unsupported CLI Commands

This appendix lists the unsupported command-line interface (CLI) commands that are displayed when you enter the question mark (?) at the switch prompt. The unsupported commands are listed by software feature and command mode.

Access Control Lists

Unsupported Privileged EXEC Commands

access-enable [**host**] [**timeout** *minutes*]

access-template [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*] [**timeout** *minutes*]

clear access-template [*access-list-number* | *name*] [*dynamic-name*] [*source*] [*destination*].

ARP Commands

Unsupported Global Configuration Commands

arp *ip-address hardware-address* **smds**

arp *ip-address hardware-address* **srp-a**

arp *ip-address hardware-address* **srp-b**

Unsupported Interface Configuration Commands

arp probe

ip probe proxy

FallBack Bridging

Unsupported Privileged EXEC Commands

clear bridge [*bridge-group*] **multicast** [**router-ports** | **groups** | **counts**] [*group-address*] [*interface-unit*] [**counts**]

clear vlan statistics

show bridge [*bridge-group*] **circuit-group** [*circuit-group*] [*src-mac-address*] [*dst-mac-address*]

show bridge [*bridge-group*] **multicast** [**router-ports** | **groups**] [*group-address*]

show bridge vlan

show interfaces crb

show interfaces {**ethernet** | **fastethernet**} [*interface* | *slot/port*] **irb**

show subscriber-policy *range*

Unsupported Global Configuration Commands

bridge *bridge-group* **bitswap_l3_addresses**

bridge *bridge-group* **bridge ip**

bridge *bridge-group* **circuit-group** *circuit-group* **pause** *milliseconds*

bridge *bridge-group* **circuit-group** *circuit-group* **source-based**

bridge cmf

bridge crb

bridge *bridge-group* **domain** *domain-name*

bridge irb

bridge *bridge-group* **mac-address-table limit** *number*

bridge *bridge-group* **multicast-source**

bridge *bridge-group* **route** *protocol*

bridge *bridge-group* **subscriber policy** *policy*

subscriber-policy *policy* [[**no** | **default**] *packet* [**permit** | **deny**]]

Unsupported Interface Configuration Commands

bridge-group *bridge-group* **cbus-bridging**

bridge-group *bridge-group* **circuit-group** *circuit-number*

bridge-group *bridge-group* **input-address-list** *access-list-number*

bridge-group *bridge-group* **input-lat-service-deny** *group-list*

bridge-group *bridge-group* **input-lat-service-permit** *group-list*

bridge-group *bridge-group* **input-lsap-list** *access-list-number*

bridge-group *bridge-group* **input-pattern-list** *access-list-number*
bridge-group *bridge-group* **input-type-list** *access-list-number*
bridge-group *bridge-group* **lat-compression**
bridge-group *bridge-group* **output-address-list** *access-list-number*
bridge-group *bridge-group* **output-lat-service-deny** *group-list*
bridge-group *bridge-group* **output-lat-service-permit** *group-list*
bridge-group *bridge-group* **output-lsap-list** *access-list-number*
bridge-group *bridge-group* **output-pattern-list** *access-list-number*
bridge-group *bridge-group* **output-type-list** *access-list-number*
bridge-group *bridge-group* **sse**
bridge-group *bridge-group* **subscriber-loop-control**
bridge-group *bridge-group* **subscriber-trunk**
bridge *bridge-group* **lat-service-filtering**
frame-relay map bridge *dlci* **broadcast**
interface bvi *bridge-group*
x25 map bridge *x.121-address* **broadcast** [*options-keywords*]

HSRP

Unsupported Global Configuration Commands

interface Async
interface BVI
interface Dialer
interface Group-Async
interface Lex
interface Multilink
interface Virtual-Template
interface Virtual-Tokenring

Unsupported Interface Configuration Commands

mtu
standby mac-refresh *seconds*
standby use-bia

Interface Configuration Commands

switchport broadcast *level*

switchport multicast *level*

switchport unicast *level*

Note These commands were replaced in IOS release 12.1(8)EA1 by the **storm-control** {**broadcast** | **multicast** | **unicast**} **level** *level* [.*level*] interface configuration command.

IP Multicast Routing

Unsupported Privileged EXEC Commands

clear ip rtp header-compression [*type number*]

The **debug ip packet** command displays packets received by the switch CPU. It does not display packets that are hardware-switched.

The **debug ip mcache** command affects packets received by the switch CPU. It does not display packets that are hardware-switched.

The **debug ip mpacket** [**detail**] [*access-list-number* [*group-name-or-address*] command affects only packets received by the switch CPU. Because most multicast packets are hardware-switched, use this command only when you know that the route will forward the packet to the CPU.

debug ip pim atm

show frame-relay ip rtp header-compression [**interface** *type number*]

The **show ip mcache** command displays entries in the cache for those packets that are sent to the switch CPU. Because most multicast packets are switched in hardware without CPU involvement, you can use this command, but multicast packet information is not displayed.

The **show ip mpacket** commands are supported but are only useful for packets received at the switch CPU. If the route is hardward-switched, the command has no effect because the CPU does not receive the packet and cannot display it.

show ip pim vc [*group-address* | *name*] [*type number*]

show ip rtp header-compression [*type number*] [**detail**]

Unsupported Global Configuration Commands

ip pim accept-rp {*address* | **auto-rp**} [*group-access-list-number*]

ip pim message-interval *seconds*

Unsupported Interface Configuration Commands

frame-relay ip rtp header-compression [**active** | **passive**]

frame-relay map ip *ip-address dlci* [**broadcast**] **compress**

frame-relay map ip *ip-address dlci* **rtp header-compression** [**active** | **passive**]

ip igmp helper-address *ip-address*

ip multicast helper-map {*group-address* | **broadcast**} {*broadcast-address* | *multicast-address*} *extended-access-list-number*

ip multicast rate-limit {**in** | **out**} [**video** | **whiteboard**] [**group-list** *access-list*] [**source-list** *access-list*] *kbps*

ip multicast use-functional

ip pim minimum-vc-rate *pps*

ip pim multipoint-signalling

ip pim nbma-mode

ip pim vc-count *number*

ip rtp compression-connections *number*

ip rtp header-compression [**passive**]

IP Unicast Routing

Unsupported Privileged EXEC or User EXEC Commands

clear ip accounting [**checkpoint**]

show cef [**drop** | **not-cef-switched**]

show ip vrf [{**brief** | **detail** | **interfaces**}] [*vrf-name*] [*output-modifiers*]

show ip accounting [**checkpoint**] [**output-packets** | **access-violations**]

Unsupported Global Configuration Commands

ip accounting-list *ip-address wildcard*

ip accounting-transits *count*

ip cef accounting [**per-prefix**] [**non-recursive**]

ip cef traffic-statistics [**load-interval** *seconds*] [**update-rate** *seconds*]]

ip flow-aggregation

ip flow-cache

ip flow-export

ip gratituitous-arps

ip local

ip reflexive-list

ip vrf

router bgp

router egp

router-isis

router iso-igrp

router mobile

router odr

router static

Unsupported Interface Configuration Commands

ip accounting

ip load-sharing [**per-packet**]

ip mtu *bytes*

ip route-cache

ip verify

ip vrf

ip unnumbered *type number*

All **ip security** commands

Unsupported VPN Configuration Commands

All

Unsupported VRF Configuration Commands

All

Unsupported Route Map Commands

set automatic-tag

set ip destination *ip-address mask*

set ip next-hop *ip-address* [*...ip-address*]

set ip next-hop *peer-address*

set ip precedence *value*

set ip *qos-group*

set tag *tag-value*

MSDP

Unsupported Privileged EXEC Commands

show access-expression

show exception

show location

show pm LINE

show smf [*interface-id*]

show subscriber-policy [*policy-number*]

show template [*template-name*]

Unsupported Global Configuration Commands

ip msdp default-peer *ip-address* | *name* [**prefix-list** *list*] (Because BGP/MBGP is not supported, use the **ip msdp peer** command instead of this command.)

RADIUS

Unsupported Global Configuration Commands

aaa nas port extended

radius-server attribute nas-port

radius-server configure

radius-server extended-portnames

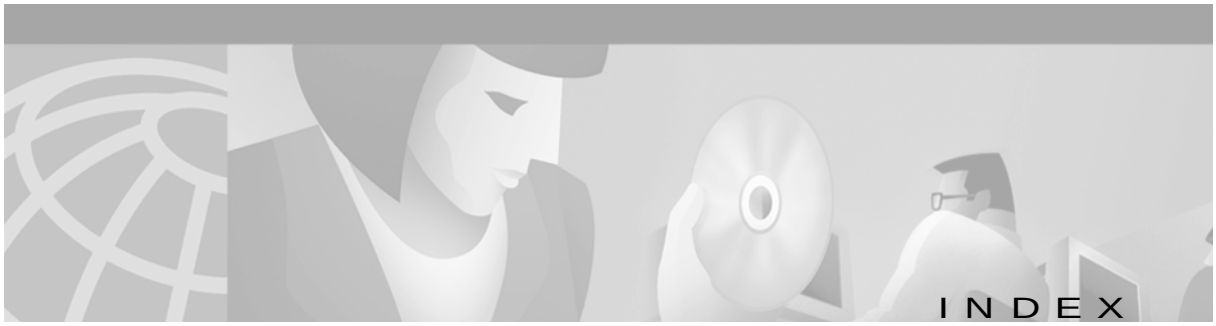

Numerics

802.1Q
  and trunk ports 8-3
  configuration limitations 9-24
  encapsulation 9-22, 9-24
  native VLAN for untagged traffic 9-29
802.1Q trunk mode 3-9
802.1X
  See port-based authentication
802.3Z flow control 8-16

A

abbreviating commands 2-3
ABRs 22-35
AC (command switch) 5-12, 5-22
access-class command 19-19
access control entries
  See ACEs
access-denied response, VMPS 9-33
access groups
  IP 19-20
  Layer 3 19-20
  viewing 19-21
accessing CMS 3-30
access levels, CMS 3-31
access lists
  See ACLs
access modes, CMS 3-31
access ports
  defined 8-2
  in switch clusters 5-11
accounting with RADIUS 6-27
accounting with TACACS+ 6-11, 6-17
ACEs
  and QoS 20-7
  defined 19-2
  Ethernet 19-2
  IP 19-2
  Layer 3 parameters 19-10
  Layer 4 parameters 19-10
ACLs
  ACEs 19-2
  any keyword 19-12
  applying
    on bridged packets 19-40
    on multicast packets 19-42
    on routed packets 19-41
    on switched packets 19-39
    time ranges to 19-15
    to an interface 19-18
    to QoS 20-7
  classifying traffic for QoS 20-27
  comments in 19-18
  compiling 19-22
  configuring with VLAN maps 19-36
  defined 19-1, 19-6
  determining fit in hardware 19-37
  displaying interface 19-22
  examples of 19-22, 20-27
  extended IP
    configuring for QoS classification 20-28
    creating 19-9
    matching criteria 19-7
  hardware and software handling 19-5

ACLs (continued)

- host keyword 19-12
- IP
  - applying to interface 19-18
  - creating 19-6
  - fragments and QoS guidelines 20-20
  - implicit deny 19-9, 19-13, 19-15
  - implicit masks 19-9
  - matching criteria 19-3, 19-7
  - named 19-14
  - options and QoS guidelines 20-20
  - undefined 19-20
  - violations, logging 19-15
  - virtual terminal lines, setting on 19-18
- limiting actions 19-37
- logging messages 19-9
- log keyword 19-15
- MAC extended 19-28, 20-29
- matching 19-6, 19-20
- monitoring 19-20
- named 19-14
- number per QoS class map 20-20
- numbers 19-7
- policy maps and QoS classification 20-20
- protocol parameters 19-10
- QoS 20-7, 20-27
- router 19-2
- standard IP
  - configuring for QoS classification 20-27
  - creating 19-8
  - matching criteria 19-7
- supported features 19-5
- support for 1-3
- time ranges 19-15
- unsupported features 19-6
- using router ACLs with VLAN maps 19-36

ACLs (continued)

- VLAN maps
  - configuration guidelines 19-28
  - configuring 19-27
  - defined 19-3

active router 23-1

addresses

- displaying the MAC address table 6-57
- dynamic
  - accelerated aging 10-10
  - changing the aging time 6-53
  - default aging 10-10
  - defined 6-51
  - learning 6-52
  - preventing frame forwarding 26-5
  - removing 6-54
- filtering frames by MAC address 26-6
- multicast
  - group address range 24-1
  - STP address management 10-8
- static
  - adding and removing 6-56
  - defined 6-51

address resolution 22-10

Address Resolution Protocol

- See ARP

adjacency tables, with CEF 22-53

administrative distances

- defined 22-62
- OSPF 22-41
- routing protocol defaults 22-56

advertisements

- CDP 13-1
- IGRP 22-30
- RIP 22-25
- VTP 9-4, 9-5, 9-25

aggregated ports

- See EtherChannel

aggregate policers 20-37

aggregate policing 1-4
aging, accelerating 10-10
aging time
  accelerated for STP 10-10, 10-29
  bridge table for fallback bridging 26-6
  MAC address table 6-53
  maximum for STP 10-18, 10-30
alarms, RMON 16-3
allowed-VLAN list 9-27
alternate routes, IGRP 22-31
Apply button 3-29
area border routers
  See ABRs
ARP
  configuring 22-11
  defined 22-10
  encapsulation 22-12
  static cache configuration 22-11
  support for 1-2
ASBRs 22-35
attributes, RADIUS
  vendor-proprietary 6-29
  vendor-specific 6-28
audience xxvii
authentication
  EIGRP 22-50
  HSRP 23-8
  local mode with AAA 6-31
  NTP associations 6-35
  RADIUS
    key 6-21
    login 6-23
  See also port-based authentication
  TACACS+
    defined 6-11
    key 6-13
    login 6-14
authentication keys, and routing protocols 22-63
authoritative time source, described 6-32
authorization with RADIUS 6-26
authorization with TACACS+ 6-11, 6-16
authorized ports with 802.1X 7-4
autoconfiguration 4-3
automatic discovery
  adding member switches 5-20
  considerations
    beyond a non-candidate device 5-8, 5-9
    brand new switches 5-11
    connectivity 5-4
    different VLANs 5-7
    management VLANs 5-8, 5-9
    non-CDP-capable devices 5-6
    non-cluster-capable devices 5-6
    routed ports 5-10
  creating a cluster standby group 5-22
  in switch clusters 5-4
  See also CDP
automatic recovery, clusters 5-12
  See also HSRP
autonegotiation
  duplex mode 1-2
  interface configuration guidelines 8-14
  mismatches 27-10
autonomous system boundary routers
  See ASBRs
Auto-RP, described 24-8
autosensing, port speed 1-2

B

BackboneFast 1-3, 10-18, 10-36
bandwidth, graphs 3-8
bandwidth for QoS
  allocating 20-50
  described 20-13

banners
configuring
login 6-51
message-of-the-day login 6-50
default configuration 6-49
when displayed 6-49
binding cluster group and HSRP group 23-9
blocking packets 12-6
booting
boot loader, function of 4-2
boot process 4-1
manually 4-13
specific image 4-14
boot loader
accessing 4-15
described 4-2
environment variables 4-15
prompt 4-15
trap-door mechanism 4-2
bootstrap router (BSR), described 24-8
BPDU guard, STP 1-3, 10-11, 10-33
bridged packets, ACLs on 19-40
bridge groups
See fallback bridging
Bridge Protocol Data Unit
See BPDU
broadcast flooding 22-20
broadcast packets
directed 22-17
flooded 22-17
broadcast storm-control command 12-3
broadcast storms 12-1, 22-17
browser configuration 3-1, 5-1
buttons, CMS 3-29

C

cables, monitoring for unidirectional links 14-1
Cancel button 3-29
candidate switch
adding 5-20
automatic discovery 5-4
defined 5-3
HC 5-22
passwords 5-20
requirements 5-3
standby group 5-22
See also command switch, cluster standby group, and member switch
caution, described xxx
CC (command switch) 5-22
CDP
automatic discovery in switch clusters 5-4
configuring 13-2
default configuration 13-2
described 13-1
disabling for routing device 13-3, 13-4
enabling and disabling on an interface 13-4
enabling and disabling on a switch 13-3
monitoring 13-5
overview 13-1
support for 1-2
transmission timer and holdtime, setting 13-2
updates 13-2
CEF 22-53
CGMP
as IGMP snooping learning method 11-6
clearing cached group entries 24-58
enabling server support 24-38
FastLeave feature 24-13
hosts
joining a group 24-12
leaving a group 24-13
joining multicast group 11-2
overview 24-11
server support only 24-11
switch support of 1-2
change notification, CMS 3-32

Cisco Discovery Protocol
See CDP
Cisco Express Forwarding
See CEF
Cisco Group Management Protocol
See CGMP
Cisco Technical Assistance Center xxxiii
CiscoWorks 2000 1-6, 18-3
classless routing 22-8
class maps for QoS
configuring 20-30
described 20-7
displaying 20-56
class of service
See CoS
clearing interfaces 8-20
CLI, managing clusters in 5-25
client mode, VTP 9-4
clock
See system clock
Cluster Management Suite
See CMS
clusters, switch
accessing 5-15
adding member switches 5-20
automatic discovery 5-4
automatic recovery 5-12
benefits 1-6
command switch configuration 5-19
compatibility 5-4
creating 5-18
creating a cluster standby group 5-22
described 5-1
LRE profile considerations 5-17
managing
through CLI 5-25
through SNMP 5-26

clusters, switch (continued)
planning considerations
automatic discovery 5-4
automatic recovery 5-12
CLI 5-25
described 5-4
host names 5-16
IP addresses 5-15
LRE profiles 5-17
passwords 5-16
SNMP 5-16, 5-26
switch-specific features 5-18
TACACS+ 5-17
redundancy 5-22
See also candidate switch, command switch, cluster standby group, member switch, and standby command switch
troubleshooting 5-24
verifying 5-24
cluster standby group
automatic recovery 5-13
considerations 5-14
creating 5-22
defined 5-2
requirements 5-3
virtual IP address 5-13
See also HSRP
cluster standby-group and HSRP group 23-9
cluster tree
described 3-5
icons 3-5
CMS
accessing 3-30
access levels 3-31
benefits 1-6
change notification 3-32
cluster tree 3-5
described 1-1, 1-5, 3-1
different versions of 3-33

CMS (continued)
displaying system messages 3-19
error checking 3-32
features 3-2
Front Panel images 3-6
Front Panel view 3-4
interaction modes 3-26
menu bar 3-15
privilege level 3-31
requirements 3-30
saving configuration changes 3-32
toolbar 3-21
tool tips 3-27
Topology view 3-10
verifying configuration changes 3-32
window components 3-28
wizards 3-26
CMS online help 3-27
Collapse Cluster view 3-11
command-line interface
abbreviating commands 2-3
command modes 2-1
described 1-5
editing features
enabling and disabling 2-6
keystroke editing 2-6
wrapped lines 2-8
error messages 2-4
filtering command output 2-8
getting help 2-3
history
changing the buffer size 2-5
described 2-5
disabling 2-5
recalling commands 2-5
no and default forms of commands 2-4
command modes 2-1
commands
abbreviating 2-3
no and default 2-4
setting privilege levels 6-8
command switch
active (AC) 5-12, 5-22
command switch with HSRP disabled (CC) 5-22
configuration conflicts 27-10
defined 5-2
enabling 5-19
passive (PC) 5-12, 5-22
password privilege levels 5-25
priority 5-12
recovery
from failure 27-7
from lost member connectivity 27-10
recovery from command-switch failure 5-12
redundant 5-12, 5-22
replacing
with another switch 27-9
with cluster member 27-7
requirements 5-2
See also candidate switch, cluster standby group, member switch, and standby command switch
standby (SC) 5-12, 5-22
community strings
configuring 5-16, 18-5
for cluster switches 18-3
in clusters 5-16
overview 18-3
SNMP 5-16
config.text 4-12
configuration changes, saving 3-32
configuration conflicts, recovering from lost member connectivity 27-10
configuration examples, network 1-7

configuration files
clearing the startup configuration B-19
creating using a text editor B-10
default name 4-12
deleting a stored configuration B-19
described B-8
downloading
automatically 4-12
preparing B-10, B-13, B-16
reasons for B-8
using FTP B-13
using RCP B-17
using TFTP B-11
guidelines for creating and using B-9
invalid combinations when copying B-5
limiting TFTP server access 18-9
obtaining with DHCP 4-7
password recovery disable considerations 6-5
specifying the filename 4-13
system contact and location information 18-9
types and location B-9
uploading
preparing B-10, B-13, B-16
reasons for B-8
using FTP B-15
using RCP B-18
using TFTP B-12
VMPS database 9-34
configuration settings, saving 4-10
configure terminal command 8-7
conflicts, configuration 27-10
congestion-avoidance techniques 20-13
congestion-management techniques 20-13, 20-15
connectivity problems 27-11
consistency checks in VTP version 2 9-6
conventions
command xxx
for examples xxx
publication xxx
text xxx
CoS 1-4, 20-2
CoS-to-DSCP map for QoS 20-39
CoS-to-egress-queue map 20-45
counters, clearing interface 8-20
Cpu q, in show forward command output 27-16
crashinfo file 27-17
cross-stack UplinkFast, STP
configuring 10-35
connecting stack ports 10-16
described 10-13
fast-convergence events 10-15
Fast Uplink Transition Protocol 10-14
limitations 10-16
normal-convergence events 10-15
Stack Membership Discovery Protocol 10-14
support for 1-3

D

database, VTP 9-15, 9-17
daylight saving time 6-44
debugging
enabling all system diagnostics 27-15
enabling for a specific feature 27-14
redirecting error message output 27-15
using commands 27-14
default commands 2-4
default configuration
802.1X 7-6
banners 6-49
booting 4-12
CDP 13-2
DNS 6-48
EIGRP 22-47

default configuration (continued)
EtherChannel 21-7
fallback bridging 26-3
HSRP 23-3
IGMP 24-31
IGMP filtering 11-20
IGMP snooping 11-5
IGRP 22-30
initial switch information 4-3
IP addressing, IP routing 22-4
IP multicast routing 24-13
Layer 2 interfaces 8-13
MAC address table 6-53
MSDP 25-4
MVR 11-15
NTP 6-35
OSPF 22-36
password and privilege level 6-3
QoS 20-18
RADIUS 6-20
RIP 22-25
RMON 16-3
SNMP 18-4
SPAN 15-7
STP 10-21
system message logging 17-3
system name and prompt 6-46
TACACS+ 6-13
UDLD 14-3
VLAN, Layer 2 Ethernet interfaces 9-24
VLANs 9-15
VMPS 9-37
VTP 9-8
default gateway 4-10, 22-15
default networks 22-56
default routes 22-56
default routing 22-2
deleting VLAN from database 9-18
description command 8-17
designing your network, examples 1-7
destination addresses, in ACLs 19-11
device discovery protocol 13-1
device icons
Front Panel view 3-5
Topology view 3-12
device labels 3-13
Device Manager 3-2
See also Switch Manager
device pop-up menu
Front Panel view 3-22
Topology view 3-24
DHCP
autoconfiguration example 4-8
client request message exchange 4-4
configuring
client side 4-3
DNS 4-6
relay device 4-6
server-side 4-5
TFTP server 4-5
lease options
for IP address information 4-5
for receiving the configuration file 4-5
overview 4-3
relationship to BOOTP 4-3
support for 1-2
DHCP relay 1-5
Differentiated Services architecture, QoS 20-1
Differentiated Services Code Point 20-2
Diffusing Update Algorithm (DUAL) 22-46
directed unicast requests 1-2
directories
changing B-3
creating and removing B-4
displaying the working B-3
discovery, clusters
See automatic discovery
display options, Topology view 3-14

Disqualification Code option 3-25
Distance Vector Multicast Routing Protocol
  See DVMRP
distance-vector protocols 22-2
distribute-list command 22-62
DNS
  and DHCP-based autoconfiguration 4-6
  default configuration 6-48
  displaying the configuration 6-49
  overview 6-47
  setting up 6-48
  support for 1-2
documentation
  feedback xxxii
  obtaining
    CD-ROM xxxi
    world wide web xxxi
  ordering xxxii
  related xxxi
document conventions xxx
domain names
  DNS 6-47
  VTP 9-8
Domain Name System
  See DNS
downloading
  configuration files
    preparing B-10, B-13, B-16
    reasons for B-8
    using FTP B-13
    using RCP B-17
    using TFTP B-11
  image files
    deleting old image B-23
    preparing B-22, B-25, B-29
    reasons for B-19
    using FTP B-26
    using RCP B-30
    using TFTP B-22
DSCP 1-4, 20-2
DSCP-to-CoS map for QoS 20-42
DSCP-to-DSCP-mutation map for QoS 20-43
DSCP-to-threshold map for QoS 20-47
DTP 1-3, 9-22
DUAL finite state machine, EIGRP 22-46
duplex mode, configuring 8-14
DVMRP
  all-DVMRP-routers multicast group address 24-11
  autosummarization
    configuring a summary address 24-54
    disabling 24-56
  connecting PIM domain to DVMRP router 24-46
  enabling unicast routing 24-50
  interoperability
    with Cisco devices 24-44
    with IOS software 24-11
  mrinfo requests, responding to 24-49
  neighbors
    advertising the default route to 24-48
    discovery with Probe messages 24-11, 24-44
    displaying information 24-49
    prevent peering with nonpruning 24-52
    rejecting nonpruning 24-51
  overview 24-11
  routes
    adding a metric offset 24-56
    advertising all 24-56
    advertising the default route to neighbors 24-48
    caching DVMRP routes learned in report messages 24-50
    changing the threshold for syslog messages 24-54
    deleting 24-58
    displaying 24-58
    favoring one over another 24-56
    limiting the number injected into MBONE 24-53
    limiting unicast route advertisements 24-44
  route table, building 24-11
  source distribution tree, building 24-11

DVMRP (continued)
support for 1-5
tunnels
configuring 24-46
displaying neighbor information 24-49
dynamic access mode 3-9
dynamic access ports
characteristics 9-3
configuring 9-38
defined 8-2
limit on number of hosts 9-40
dynamic addresses
See addresses
Dynamic Host Configuration Protocol
See DHCP
dynamic port VLAN membership
configuration example 9-40
overview 9-34
reconfirming 9-39
troubleshooting 9-40
types of connections 9-38
VMPS database configuration file 9-34
dynamic routing 22-2
Dynamic Trunking Protocol
See DTP

E

editing features
enabling and disabling 2-6
keystrokes used 2-6
wrapped lines 2-8
Egress q, in show forward command output 27-16
EIGRP
and IGRP 22-48
authentication 22-50
components 22-46
configuring 22-48
default configuration 22-47
EIGRP (continued)
definition 22-46
interface parameters, configuring 22-49
monitoring 22-51
support for 1-4
enable password 6-4
enable secret password 6-4
encapsulation types, Ethernet trunk 9-24
encryption for passwords 6-4
Enhanced IGRP
See EIGRP
environment variables
function of 4-16
location in Flash 4-15
equal-cost routing 1-4, 22-54
error checking, CMS 3-32
error messages
during command entry 2-4
setting the display destination device 17-4
severity levels 17-8
system message format 17-2
EtherChannel
automatic creation of 21-3
channel groups
binding physical and logical interfaces 21-2
numbering of 21-2
configuration guidelines 21-8
configuring
Layer 2 interfaces 21-9
Layer 3 physical interfaces 21-12
Layer 3 port-channel logical interfaces 21-11
default configuration 21-7
destination MAC address forwarding 21-5
detecting misconfigurations 10-37
displaying status 21-16
forwarding methods 21-5, 21-13
interaction
with STP 21-8
with VLANs 21-8

EtherChannel (continued)
Layer 3 interface 22-3
load balancing 21-5, 21-13
logical interfaces, described 21-2
number of interfaces per 21-1
overview 21-1
PAgP
aggregate-port learners 21-5
compatibility with Catalyst 1900 21-14
displaying status 21-16
interaction with other features 21-5
learn method and priority configuration 21-14
modes 21-4
overview 21-3
silent mode 21-4
support for 1-2
port-channel interfaces
described 21-2
numbering of 21-2
port groups 8-3
source MAC address forwarding 21-5
support for 1-2
Ethernet VLAN
defaults and ranges 9-16
modifying 9-18
events, RMON 16-3
examples
conventions for xxx
network configuration 1-7
Expand Cluster view 3-11
expedite queue for QoS
10/100 Ethernet ports
allocating bandwidth 20-54
configuring 20-54
described 20-15

expedite queue for QoS (continued)
Gigabit-capable Ethernet ports
allocating bandwidth 20-50
configuring 20-50
described 20-13
expert mode 3-26
extended system ID for STP 10-3, 10-23
Extensible Authentication Protocol over LAN 7-1
exterior routes, IGRP 22-30

F

fallback bridging
and protected ports 26-4
bridge groups
creating 26-4
described 26-2
displaying 26-12
function of 26-2
number supported 26-4
removing 26-4
bridge table
changing the aging time 26-6
clearing 26-12
displaying 26-12
connecting interfaces with 8-6
default configuration 26-3
described 26-1
frame forwarding
filtering by MAC address 26-6
flooding packets 26-2
for static addresses 26-5
forwarding packets 26-2
preventing for dynamically learned stations 26-5
to static addresses 26-5
overview 26-1

fallback bridging (continued)
STP
disabling on an interface 26-12
forward-delay interval 26-10
hello BPDU interval 26-10
interface priority 26-8
maximum-idle interval 26-11
path cost 26-9
switch priority 26-8
VLAN-bridge STP 26-1, 26-2
support for 1-4
SVIs and routed ports 26-1
VLAN-bridge STP 10-9
fallback VLAN name 9-34
fan fault indication 3-5
Fast Uplink Transition Protocol 10-14
feedback to Cisco Systems, web xxxii
FIB 22-53
fiber-optic, detecting unidirectional links 14-1
files
copying B-4
crashinfo
description 27-17
displaying the contents of 27-17
location 27-17
deleting B-5
displaying the contents of B-8
tar
creating B-6
displaying the contents of B-6
extracting B-7
image file format B-20
file system
displaying available file systems B-2
displaying file information B-3
local file system names B-1
network file system names B-4
setting the default B-3
filtering
in a VLAN 19-27
non-IP traffic 19-28
show and more command output 2-8
with fallback bridging 26-6
filters, IP
See ACLs, IP
Flash device, number of B-1
flash updates, IGRP 22-31
flooded traffic, blocking 12-6
flow-based packet classification 1-4
flowcharts
QoS classification 20-6
QoS policing and marking 20-10
QoS queueing and scheduling
10/100 ports 20-15
Gigabit-capable ports 20-12
flow control 1-2, 8-16
forward-delay time, STP 10-7, 10-29
Forwarding Information Base
See FIB
forwarding non-routable protocols 26-1
Front Panel images, CMS 3-6
Front Panel view
cluster tree 3-5
described 3-4
pop-up menus 3-22
port icons 3-6
port LEDs 3-8
RPS LED 3-7
switch images 3-6
FTP
accessing MIB files A-2
configuration files
downloading B-13
overview B-12
preparing the server B-13
uploading B-15

FTP (continued)
image files
deleting old image B-28
downloading B-26
preparing the server B-25
uploading B-28

G

GBICs
1000BASE-LX/LH module 1-9
1000BASE-SX module 1-9
1000BASE-T module 1-9
1000BASE-ZX module 1-9
GigaStack 1-9
get-bulk-request operation 18-2
get-next-request operation 18-2, 18-3
get-request operation 18-2, 18-3
get-response operation 18-2
Gigabit GBIC modules
See GBICs
Gigabit Interface Converters
See GBICs
GigaStack GBIC
fast transition of redundant link 10-13
See also GBICs
global configuration mode 2-2
graphs, bandwidth 3-8
guide
audience xxvii
organization of xxviii
purpose of xxvii
guide mode 1-6, 3-26

H

hardware, determining ACL configuration fit 19-37
HC (candidate switch) 5-22
hello time, STP 10-29
help, for the command line 2-3
Help button, CMS 3-29
Help Contents 3-27
history
changing the buffer size 2-5
described 2-5
disabling 2-5
recalling commands 2-5
history table, level and number of syslog messages 17-10
host name list, CMS 3-28
host names
abbreviations appended to 5-22
in clusters 5-16
hosts, limit on dynamic ports 9-40
Hot Standby Router Protocol
See HSRP
HP OpenView 1-6
HSRP
authentication string 23-8
automatic cluster recovery 5-13
binding to cluster group 23-9
cluster standby group considerations 5-14
command-switch redundancy 1-1
configuring 23-3
default configuration 23-3
definition 23-1
monitoring 23-10
overview 23-1
priority 23-6
routing redundancy 1-4
timers 23-8
tracking 23-6
See also clusters, cluster standby group, and standby command switch
HTTP access 3-30, 3-31
Hypertext Transfer Protocol
See HTTP access

I

ICMP
redirect messages 22-15
support for 1-5
time exceeded messages 27-12
traceroute and 27-12
unreachable messages 19-5
unreachables and ACLs 19-6
ICMP ping
executing 27-11
overview 27-11
ICMP Router Discovery Protocol
See IRDP
icons
cluster tree 3-5
colors
cluster tree 3-5
Topology view 3-14
editable table cell 3-29
Front Panel view 3-6
multilink 3-23
sorting 3-29
toolbar 3-21
Topology view 3-12
web link 3-29
IFS 1-2
IGMP
configuring the switch
as a member of a group 24-34
statically connected member 24-36
controlling access to groups 24-35
default configuration 24-31
deleting cache entries 24-58
displaying groups 24-58
fast switching 24-37
host-query interval, modifying 24-36
joining multicast group 11-2
join messages 11-2

IGMP (continued)
leave processing, enabling 11-9
leaving multicast group 11-4
multicast reachability 24-34
overview 24-3
queries 11-3
support for 1-2
Version 1
changing to Version 2 24-32
hosts joining a group 24-3
hosts leaving a group 24-3
membership queries 24-3
overview 24-3
query-response model 24-3
Version 2
changing to Version 1 24-32
enhancements over Version 1 24-4
hosts leaving a group 24-4
maximum query response time value 24-33
new features 24-4
overview 24-4
pruning groups 24-33
query timeout value 24-32
IGMP filtering
configuring 11-20
default configuration 11-20
described 11-20
monitoring 11-24
IGMP groups, setting the maximum number 11-23
IGMP profile
applying 11-22
configuration mode 11-20
configuring 11-21
IGMP snooping
configuring 11-5
default configuration 11-5
definition 11-1
enabling and disabling 11-5
global configuration 11-5

IGMP snooping (continued)
Immediate Leave 11-4
method 11-6
monitoring 11-9
support for 1-2
VLAN configuration 11-6
IGP 22-35
IGRP
advertisements 22-30
alternate routes 22-31
configuring 22-32
default configuration 22-30
described 22-30
exterior routes 22-30
flash updates 22-31
interior routes 22-30
load balancing 22-31
poison-reverse updates 22-31
split horizon 22-34
support for 1-4
system routes 22-30
traffic sharing 22-32
unequal-cost load balancing 22-31
Immediate-Leave, IGMP 11-4
interaction modes, CMS 3-26
interface
number 8-7
range macros 8-11
interface command 8-7
interface configuration mode 2-2
interfaces
configuration guidelines 8-14
configuring 8-7
configuring duplex mode 8-14
configuring speed 8-14
counters, clearing 8-20
described 8-17
descriptive name, adding 8-17
displaying information about 8-18
interfaces (continued)
flow control 8-16
management 1-5
monitoring 8-18
naming 8-17
physical, identifying 8-7
range of 8-9
restarting 8-21
shutting down 8-21
supported 8-6
types of 8-1
interfaces range macro command 8-11
Interior Gateway Protocol
See IGP
Interior Gateway Routing Protocol
See IGRP
interior routes, IGRP 22-30
Internet Control Message Protocol
See ICMP
Internet Group Management Protocol
See IGMP
Inter-Switch Link
See ISL
inter-VLAN routing 1-4, 22-2
inventory, cluster 5-24
IOS File System
See IFS
ip access group command 19-20
IP ACLs
applying to an interface 19-18
extended, creating 19-9
for QoS classification 20-7
implicit deny 19-9, 19-13, 19-15
implicit masks 19-9
logging 19-15
named 19-14
standard, creating 19-8
undefined 19-20
virtual terminal lines, setting on 19-18

IP addresses
candidate or member 5-3, 5-15
classes of 22-5
cluster access 5-2
command switch 5-2, 5-13, 5-15
default configuration 22-4
for IP routing 22-4
MAC address association 22-10
monitoring 22-21
redundant clusters 5-13
standby command switch 5-13, 5-15
See also IP information
IP broadcast address 22-20
ip cef command 22-53
IP directed broadcasts 22-17
ip igmp profile command 11-20
IP information
assigned
manually 4-10
through DHCP-based autoconfiguration 4-3
default configuration 4-3
IP multicast routing
addresses
all-hosts 24-1
all-multicast-routers 24-1
all-PIM-routers 24-10
Cisco-RP-Announce 24-8
Cisco-RP-Discovery 24-8
host group address range 24-1
administratively-scoped boundaries, described 24-42
and IGMP snooping 11-1, 11-5

IP multicast routing (continued)
Auto-RP
adding to an existing sparse-mode cloud 24-19
benefits of 24-18
clearing the cache 24-58
configuration guidelines 24-15
IOS release 24-5
overview 24-8
preventing candidate RP spoofing 24-21
preventing join messages to false RPs 24-20
setting up in a new internetwork 24-19
using with BSR 24-27
bootstrap router
configuration guidelines 24-15
configuring candidate BSRs 24-25
configuring candidate RPs 24-26
defining the IP multicast boundary 24-24
defining the PIM domain border 24-22
IOS release 24-5
overview 24-8
using with Auto-RP 24-27
Cisco implementation 24-2
configuring
basic multicast routing 24-15
IP multicast boundary 24-42
TTL threshold 24-40
default configuration 24-13
enabling
multicast forwarding 24-15
PIM mode 24-16
group-to-RP mappings
Auto-RP 24-8
BSR 24-8

IP multicast routing (continued)
MBONE
deleting sdr cache entries 24-58
described 24-39
displaying sdr cache 24-59
enabling sdr listener support 24-39
limiting DVMRP routes advertised 24-53
limiting sdr cache entry lifetime 24-39
SAP packets for conference session announcement 24-39
Session Directory (sdr) tool, described 24-39
monitoring
packet rate loss 24-59
peering devices 24-59
tracing a path 24-59
multicast forwarding, described 24-9
PIMv1 and PIMv2 interoperability 24-14
protocol interaction 24-2
reverse path check (RPF) 24-9
routing table
deleting 24-58
displaying 24-58
RP
assigning manually 24-17
configuring Auto-RP 24-18
configuring PIMv2 BSR 24-22
monitoring mapping information 24-27
using Auto-RP and BSR 24-27
statistics, displaying system and network 24-58
TTL thresholds, described 24-40
See also CGMP
See also DVMRP
See also IGMP
See also PIM
IP precedence 20-2
IP-precedence-to-DSCP map for QoS 20-40
IP protocols
in ACLs 19-11
routing 1-4
IP routes, monitoring 22-64
IP routing
connecting interfaces with 8-6
enabling 22-24
IP traceroute
executing 27-13
overview 27-12
IP unicast routing
address resolution 22-10
administrative distances 22-56, 22-62
ARP 22-10
assigning IP addresses to Layer 3 interfaces 22-6
authentication keys 22-63
broadcast address 22-20
broadcast flooding 22-20
broadcast packets 22-17
broadcast storms 22-17
classless routing 22-8
configuring static routes 22-55
default addressing configuration 22-4
default gateways 22-15
default networks 22-56
default routes 22-56
default routing 22-2
directed broadcasts 22-17
dynamic routing 22-2
enabling 22-24
EtherChannel Layer 3 interface 22-3
IGP 22-35
inter-VLAN 22-2
IP addressing
classes 22-5
configuring 22-4
IRDP 22-15
Layer 3 interfaces 22-3
MAC address and IP address 22-10
passive interfaces 22-61

IP unicast routing (continued)
protocols
distance-vector 22-2
dynamic 22-2
link-state 22-2
proxy ARP 22-10
redistribution 22-57
reverse address resolution 22-10
routed ports 22-3
static routing 22-2
steps to configure 22-3
subnet mask 22-5
subnet zero 22-8
supernet 22-8
UDP 22-19
with SVIs 22-3
See also EIGRP
See also IGRP
See also OSPF
See also RIP
ip unreachables command 19-5
IRDP
configuring 22-16
definition 22-15
support for 1-5
ISL
and trunk ports 8-3
encapsulation 1-3, 9-22, 9-24
ISL trunk mode 3-9

J

Java plug-in configuration 3-1, 5-1
join messages, IGMP 11-2

L

Layer 2 frames, classification with CoS 20-2
Layer 2 interfaces, default configuration 8-13
Layer 2 trunks 9-22
Layer 3 features 1-4
Layer 3 interfaces
assigning IP addresses to 22-6
changing from Layer 2 mode 22-6
types of 22-3
Layer 3 packets, classification methods 20-2
Layer 3 parameters of ACEs 19-10
Layer 4 parameters of ACEs 19-10
leave processing, IGMP 11-9
LEDs
port 3-8
port modes 3-8
RPS 3-7
legend, CMS icons and labels 3-20
line configuration mode 2-2
link icons, Topology view 3-13
link labels 3-13
link pop-up menu, Topology view 3-23
links, unidirectional 14-1
link state advertisements (LSAs) 22-39
link-state protocols 22-2
lists, CMS 3-29
load balancing, IGRP 22-31
logging messages, ACL 19-9
login authentication with RADIUS 6-23
login authentication with TACACS+ 6-14
login banners 6-49
log messages
See system message logging
Long-Reach Ethernet (LRE) technology 1-8

M

MAC addresses
  aging time 6-53
  and VLAN association 6-52
  building the address table 6-52
  default configuration 6-53
  displaying 6-57
  dynamic
    learning 6-52
    removing 6-54
  in ACLs 19-28
  IP address association 22-10
  static
    adding 6-56
    characteristics of 6-56
    removing 6-56
MAC address multicast entries, monitoring 11-10
MAC address-to-VLAN mapping 9-33
MAC extended access lists 19-28, 20-4, 20-29
manageability features 1-2
management
  interfaces 1-5
  options 1-5
management options, CMS 3-1
management VLAN
  considerations in clusters 5-8, 5-9
mapping tables for QoS
  configuring
    CoS-to-DSCP 20-39
    CoS-to-egress-queue 20-45
    DSCP 20-39
    DSCP-to-CoS 20-42
    DSCP-to-DSCP-mutation 20-43
    DSCP-to-threshold 20-47
    IP-precedence-to-DSCP 20-40
    policed-DSCP 20-41
  described 20-11
marking
  action in policy map 20-32
  action with aggregate policers 20-37
  described 20-3, 20-8
matching, ACLs 19-6
maximum aging time, STP 10-30
maximum-paths command 22-54
membership mode, VLAN port 3-9, 9-3
member switch
  adding 5-20
  automatic discovery 5-4
  defined 5-2
  managing 5-25
  passwords 5-15
  recovering from lost connectivity 27-10
  requirements 5-3
  See also candidate switch, cluster standby group, and standby command switch
memory, optimizing 6-57
menu bar
  described 3-15
  variations 3-15
messages
  logging ACL violations 19-15
  system 3-19
  to users through banners 6-49
metric translations, between routing protocols 22-60
MIBs
  accessing files with FTP A-2
  location of files A-2
  overview 18-1
  SNMP interaction with 18-3
  supported A-1
minimum-reserve levels
  assigning to a queue 20-15, 20-53
  configuring the buffer size 20-16, 20-53
  default size 20-15
mini-point-of-presence
  See POP

mirroring traffic for analysis 15-1
mismatches, autonegotiation 27-10
Mode button 3-8
modes
  access to CMS 3-31
  port 3-8
  VLAN port membership 3-9
Modify button 3-29
modules, GBIC
  1000BASE-LX/LH 1-9
  1000BASE-SX 1-9
  1000BASE-T 1-9
  1000BASE-ZX 1-9
  GigaStack 1-9
monitoring
  access groups 19-20
  ACLs 19-20
  cables for unidirectional links 14-1
  CDP 13-5
  CEF 22-54
  EIGRP 22-51
  fallback bridging 26-12
  features 1-5
  HSRP 23-10
  IGMP filters 11-24
  IGMP snooping 11-9
  interfaces 8-18
  IP address tables 22-21
  IP multicast routing 24-57
  IP routes 22-64
  MSDP peers 25-19
  multicast router interfaces 11-10
  MVR 11-18
  network traffic for analysis with probe 15-1
  OSPF 22-44
  port blocking 12-11
  port protection 12-11
  RP mapping information 24-27
  source-active messages 25-19

monitoring (continued)
  speed and duplex mode 8-15
  traffic flowing among switches 16-1
  traffic suppression 12-11
  VLAN
    filters 19-33
    maps 19-33
  VMPS 9-40
  VTP 9-13
  VTP database 9-21
MSDP
  and dense-mode regions
    sending SA messages to 25-17
    specifying the originating address 25-18
  benefits of 25-3
  clearing MSDP connections and statistics 25-19
  controlling source information
    forwarded by switch 25-12
    originated by switch 25-8
    received by switch 25-14
  default configuration 25-4
  filtering
    incoming SA messages 25-14
    SA messages to a peer 25-12
    SA requests from a peer 25-11
  join latency, defined 25-6
  meshed groups
    configuring 25-16
    defined 25-16
  originating address, changing 25-18
  overview 25-1
  peer-RPF flooding 25-2
  peers
    configuring a default 25-4
    monitoring 25-19
    peering relationship, overview 25-1
    requesting source information from 25-8
    shutting down 25-16

MSDP (continued)
source-active messages
caching 25-6
clearing cache entries 25-19
defined 25-2
filtering from a peer 25-11
filtering incoming 25-14
filtering to a peer 25-12
limiting data with TTL 25-14
monitoring 25-19
restricting advertised sources 25-9
multicast groups
and IGMP snooping 11-5
Immediate Leave 11-4
joining 11-2
leaving 11-4
static joins 11-8
multicast packets
ACLs on 19-42
blocking 12-6
multicast router interfaces, monitoring 11-10
multicast router ports, adding 11-7
Multicast Source Discovery Protocol
See MSDP
multicast storm-control command 12-3
multicast storms 12-1
Multicast VLAN Registration
See MVR
Multilink Decomposer window 3-23
multilink icon 3-23
MVR
configuring interfaces 11-16
default configuration 11-15
described 11-12
modes 11-16
monitoring 11-18
setting global parameters 11-15
support for 1-2

N

named IP ACLs 19-14
native VLANs 9-29
negotiate trunk mode 3-9
neighbor discovery/recovery, EIGRP 22-46
neighboring devices, types of 3-12
network configuration examples
increasing network performance 1-7
large network 1-13
providing network services 1-8
small to medium-sized network 1-11
network design
performance 1-7
services 1-8
network management
configuring CDP 13-1
configuring RMON 16-1
configuring SNMP 18-1
Network Time Protocol
See NTP
no commands 2-4
non-IP traffic filtering 19-28
nontrunking mode 9-23
no switchport command 8-4
note, described xxx
not-so-stubby areas
See NSSA
NSSA, OSPF 22-39
NTP
associations
authenticating 6-35
defined 6-33
enabling broadcast messages 6-37
peer 6-36
server 6-36
default configuration 6-35
displaying the configuration 6-41
overview 6-32

NTP (continued)
restricting access
creating an access group 6-39
disabling NTP services per interface 6-40
source IP address, configuring 6-40
stratum 6-32
support for 1-2
synchronizing devices 6-36
time
services 6-33
synchronizing 6-32

O

OK button 3-29
online help 3-27
Open Shortest Path First
See OSPF
optimizing system resources 6-57
options, management 1-5
OSPF
area parameters, configuring 22-39
configuring 22-37
default configuration
metrics 22-41
route 22-41
settings 22-36
described 22-35
interface parameters, configuring 22-38
LSA group pacing 22-43
monitoring 22-44
router IDs 22-43
route summarization 22-41
support for 1-4
virtual links 22-41
out-of-profile markdown 1-4
output interface, getting information about 27-16
overheating indication, switch 3-5

P

packet modification, with QoS 20-17
PAgP
See EtherChannel
parallel links 9-29
parallel paths, in routing tables 22-54
passive interfaces
configuring 22-61
OSPF 22-42
passwords
default configuration 6-3
disabling recovery of 6-5
encrypting 6-4
for security 1-3
in clusters 5-16, 5-20
in CMS 3-30
overview 6-1
recovery of 27-3
setting
enable 6-3
enable secret 6-4
Telnet 6-6
with usernames 6-7
VTP domain 9-9
path cost, STP 10-27
PC (passive command switch) 5-12, 5-22
performance, network design 1-7
performance features 1-2
per-VLAN Spanning Tree (PVST) 1-3, 10-2
per-VLAN Spanning Tree+ (PVST+) 10-8
physical ports 8-2

PIM
default configuration 24-13
dense mode
(S,G) notation 24-6
graft messages 24-6
overview 24-5
pruning and SPT 24-5
rendezvous point (RP), described 24-7
RPF lookups 24-10
displaying neighbors 24-59
enabling a mode 24-16
neighbor discovery and adjacencies 24-10
overview 24-5
router-query message interval, modifying 24-30
shared tree and source tree, overview 24-28
shortest path tree, delaying the use of 24-29
sparse mode
(*,G) notation 24-7
join messages and shared tree 24-7
overview 24-7
prune messages 24-8
RPF lookups 24-10
support for 1-5
versions
interoperability 24-14
supported 24-5
troubleshooting interoperability problems 24-28
v2 improvements 24-5
PIM-DVMRP, as snooping method 11-6
ping
character output description 27-12
executing 27-11
overview 27-11
planning considerations, switch clusters
LRE profiles 5-17
switch-specific features 5-18
poison-reverse updates, IGRP 22-31
policed-DSCP map for QoS 20-41
policers
configuring
for each matched traffic class 20-32
for more than one traffic class 20-37
described 20-3
displaying 20-56
number of 1-4, 20-9
types of 20-8
policing
described 20-3
token bucket algorithm 20-8
policy maps for QoS
characteristics of 20-32
configuring 20-32
described 20-7
displaying 20-56
POP 1-14
Port Aggregation Protocol
See EtherChannel
See PAgP
port-based authentication
authentication server
defined 7-2
RADIUS server 7-2
client, defined 7-2
configuration guidelines 7-7
configuring
manual re-authentication of a client 7-11
quiet period 7-11
RADIUS server 7-10
RADIUS server parameters on the switch 7-9
switch-to-client frame-retransmission number 7-13
switch-to-client retransmission time 7-12
default configuration 7-6
described 7-1
device roles 7-2
displaying statistics 7-14
EAPOL-start frame 7-3
EAP-request/identity frame 7-3

port-based authentication (continued)
EAP-response/identity frame 7-3
enabling
802.1X authentication 7-8
periodic re-authentication 7-10
encapsulation 7-2
initiation and message exchange 7-3
method lists 7-8
ports
authorization state and dot1x port-control command 7-4
authorized and unauthorized 7-4
resetting to default values 7-14
switch
as proxy 7-2
RADIUS client 7-2
topologies, supported 7-4
port blocking 1-2, 12-6
port-channel
See EtherChannel
Port Fast, STP 9-36, 10-10, 10-32
Port Fast mode 1-3
port icons, Front Panel view 3-6
port LEDs 3-8
DUPLX 3-8
port modes 3-8
SPEED 3-8
STAT 3-8
port membership modes, VLAN 3-9, 9-3
port modes
described 3-8
LEDs 3-8
port pop-up menu, Front Panel view 3-22
port priority, STP 10-26
ports
802.1Q trunk 3-9
access 8-2
blocking 12-6
dynamic access 3-9, 9-3
ports (continued)
dynamic VLAN membership, reconfirming 9-39
forwarding, resuming 12-7
ISL trunk 3-9
negotiate trunk 3-9
protected 12-5
routed 8-4
secure 12-8
static-access 3-9, 9-3, 9-19
switch 8-2
trunks 9-3, 9-22
VLAN assignments 9-19
port security
configuring 12-9
default configuration 12-9
described 12-8
displaying 12-11
violations 12-8
with other features 12-9
port-shutdown response, VMPS 9-33
preferential treatment of traffic
See QoS
preventing unauthorized access 6-1
priority, HSRP 23-6
private VLAN edge ports
See protected ports
privileged EXEC mode 2-2
privilege levels
access modes
read-only 3-31
read-write 3-31
changing the default for lines 6-9
CMS 3-31
command switch 5-25
exiting 6-10
logging into 6-10
mapping on member switches 5-25
overview 6-2, 6-8
setting a command with 6-8

protected ports 1-3, 12-5
protocol-dependent modules, EIGRP 22-47
Protocol-Independent Multicast Protocol
  See PIM
proxy ARP
  configuring 22-13
  definition 22-10
  with IP routing disabled 22-14
pruning, VTP
  enabling 9-13
  enabling on a port 9-28
  examples 9-7
  overview 9-6
pruning-eligible list
  changing 9-28
  VLANs 9-13
publications, related xxxi
PVST 1-3

Q

QoS
  basic model 20-3
  classification
    class maps, described 20-7
    defined 20-3
    flowchart 20-6
    forwarding treatment 20-3
    in frames and packets 20-2
    IP ACLs, described 20-5, 20-7
    MAC ACLs, described 20-4, 20-7
    policy maps, described 20-7
    port default, described 20-4
    trust DSCP, described 20-4
    trusted CoS, described 20-4
    trust IP precedence, described 20-4
    types for IP traffic 20-5
    types for non-IP traffic 20-4

QoS (continued)

  class maps
    configuring 20-30
    displaying 20-56
  configuration examples
    common wiring closet 20-57
    distribution layer 20-59
    intelligent wiring closet 20-58
  configuration guidelines 20-20
  configuring
    aggregate policers 20-37
    default port CoS value 20-24
    DSCP maps 20-39
    DSCP trust states bordering another domain 20-25
    egress queues on 10/100 Ethernet ports 20-51
    egress queues on Gigabit-capable Ethernet ports 20-44
    IP extended ACLs 20-28
    IP standard ACLs 20-27
    MAC ACLs 20-29
    policy maps 20-32
    port trust states within the domain 20-22
  default configuration 20-18
  displaying statistics 20-56
  enabling globally 20-21
  flowcharts
    classification 20-6
    policing and marking 20-10
    queueing and scheduling 20-12, 20-15
  implicit deny 20-7

QoS (continued)

- mapping tables
  - CoS-to-DSCP 20-39
  - CoS-to-egress-queue 20-45
  - displaying 20-56
  - DSCP-to-CoS 20-42
  - DSCP-to-DSCP-mutation 20-43
  - DSCP-to-threshold 20-47
  - IP-precedence-to-DSCP 20-40
  - policed-DSCP 20-41
  - types of 20-11
- marked-down actions 20-35
- marking, described 20-3, 20-8
- overview 20-1
- packet modification 20-17
- policers
  - configuring 20-35, 20-37
  - described 20-8
  - displaying 20-56
  - number of 20-9
  - types of 20-8
- policies, attaching to an interface 20-9
- policing
  - described 20-3, 20-8
  - token bucket algorithm 20-8
- policy maps
  - characteristics of 20-32
  - configuring 20-32
  - displaying 20-56
- queueing, defined 20-3

QoS (continued)

- queues
  - CoS-to-egress-queue map 20-45
  - for 10/100 Ethernet ports 20-15
  - high priority (expedite) 20-13, 20-50
  - minimum-reserve levels 20-53
  - serviced by WRR 20-13, 20-16
  - size of 20-13, 20-15
  - size ratios 20-46
  - tail-drop threshold percentages 20-13, 20-47
  - WRED drop-percentage thresholds 20-13, 20-48
  - WRR scheduling 20-50
- scheduling
  - allocating bandwidth on 10/100 Ethernet ports 20-54
  - allocating bandwidth on Gigabit-capable ports 20-50
  - defined 20-3
- support for 1-4
- tail drop
  - configuring drop threshold percentages 20-47
  - described 20-13
- trust states 20-4
- WRED
  - configuring drop-percentage thresholds 20-48
  - described 20-14
- WRR scheduling 20-50

quality of service

- See QoS

queries, IGMP 11-3

R

RADIUS

- attributes
  - vendor-proprietary 6-29
  - vendor-specific 6-28

RADIUS (continued)
configuring
accounting 6-27
authentication 6-23
authorization 6-26
communication, global 6-21, 6-28
communication, per-server 6-20, 6-21
multiple UDP ports 6-20
default configuration 6-20
defining AAA server groups 6-24
displaying the configuration 6-30
identifying the server 6-20
limiting the services to the user 6-26
method list, defined 6-19
operation of 6-19
overview 6-18
suggested network environments 6-18
tracking services accessed by user 6-27
Random Early Detection, described 20-14
range
macro 8-11
of interfaces 8-9
RARP 22-10
rcommand command 5-25
RCP
configuration files
downloading B-17
overview B-16
preparing the server B-16
uploading B-18
image files
deleting old image B-32
downloading B-30
preparing the server B-29
uploading B-32
read-only access mode 3-31
read-write access mode 3-31
reconfirmation interval, VMPS, changing 9-39
recovery procedures 27-1
redundancy
EtherChannel 21-1
features 1-3
HSRP 23-1
STP
backbone 10-9
multidrop backbone 10-13
path cost 9-32
port priority 9-30
redundant clusters
See cluster standby group
redundant power system
See RPS
Refresh button 3-29
reliable transport protocol, EIGRP 22-46
reloading software 4-17
Remote Authentication Dial-In User Service
See RADIUS
Remote Copy Protocol
See RCP
Remote Network Monitoring
See RMON
remove vlan-list parameter 9-27
resetting a UDLD-shutdown interface 14-4
restricting access
NTP services 6-38
overview 6-1
passwords and privilege levels 6-2
RADIUS 6-17
TACACS+ 6-10
retry count, VMPS, changing 9-39
reverse address resolution 22-10
Reverse Address Resolution Protocol
See RARP
RFC
1058, RIP 22-25
1112, IP multicast and IGMP 11-2
1157, SNMPv1 18-2
1166, IP addresses 22-5

RFC (continued)
1253, OSPF 22-35
1305, NTP 6-32
1587, NSSAs 22-35
1757, RMON 16-2
1901, SNMPv2C 18-2
1902 to 1907, SNMPv2 18-2
2236, IP multicast and IGMP 11-2
RIP
advertisements 22-25
authentication 22-28
configuring 22-26
default configuration 22-25
described 22-25
hop counts 22-25
split horizon 22-29
summary addresses 22-29
support for 1-4
RMON
default configuration 16-3
displaying status 16-6
enabling alarms and events 16-3
groups supported 16-2
overview 16-1
statistics
collecting group Ethernet 16-5
collecting group history 16-5
support for 1-5
root guard, STP 1-3, 10-20, 10-36
root switch, STP 10-22
route calculation timers, OSPF 22-42
routed packets, ACLs on 19-41
routed ports
configuring 22-3
defined 8-4
in switch clusters 5-10
routed ports, IP addresses on 8-22, 22-3
router ACLs 19-2
router ID, OSPF 22-43
route summarization, OSPF 22-41
routing
default 22-2
dynamic 22-2
redistribution of information 22-57
static 22-2
Routing Information Protocol
See RIP
routing protocol administrative distances 22-56
RPS LED 3-7
running configuration, saving 4-10

S

saving changes in CMS 3-32
SC (standby command switch) 5-12, 5-22
scheduled reloads 4-17
SDM
configuring 6-59
described 6-57
templates
number of 6-57
resources used for Fast Ethernet switches 6-58
resources used for Gigabit Ethernet switches 6-58
secure ports, configuring 12-8
security, port 12-8
security features 1-3
sequence numbers in log messages 17-8
server mode, VTP 9-4
set-request operation 18-3
setup program, failed command switch replacement 27-7, 27-9
severity levels, defining in system messages 17-8
show access-lists hw-summary command 19-6
show cdp traffic command 13-5
show cluster members command 5-25
show configuration command 8-17
show fm commands 19-37
show forward command 27-15

show interfaces command 8-15, 8-17
show running-config command
displaying ACLs 19-19, 19-30, 19-32
interface description in 8-17
shutdown command on interfaces 8-21
Simple Network Management Protocol
See SNMP
SNAP 13-1
SNMP
accessing MIB variables with 18-3
agent
described 18-3
disabling 18-5
community strings
configuring 18-5
for cluster switches 18-3
overview 18-3
configuration examples 18-10
default configuration 18-4
in-band management 1-2
in clusters 5-16
limiting access by TFTP servers 18-9
limiting system log messages to NMS 17-10
manager functions 1-6, 18-2
managing clusters with 5-26
MIBs
location of A-2
supported A-1
overview 18-1, 18-3
status, displaying 18-10
system contact and location 18-9
trap manager, configuring 18-8
traps
described 18-2
enabling 18-7
enabling MAC address notification 6-54
overview 18-1, 18-3
types of 18-7
versions supported 18-2
snooping, IGMP 11-1
software images
location in Flash B-20
recovery procedures 27-2
scheduling reloads 4-17
tar file format, described B-20
See also downloading and uploading
source addresses, in ACLs 19-11
SPAN
configuration guidelines 15-7
default configuration 15-7
destination ports 15-4
displaying status 15-13
interaction with other features 15-5
monitored ports 15-4
monitoring ports 15-4
overview 1-5, 15-1
received traffic 15-3
sessions
creating 15-8
defined 15-2
limiting source traffic to specific VLANs 15-12
monitoring VLANs 15-11
removing destination (monitoring) ports 15-10
removing source (monitored) ports 15-10
specifying monitored ports 15-8
source ports 15-4
transmitted traffic 15-3
VLAN-based 15-5
Spanning Tree Protocol
See STP
speed, configuring on interfaces 8-14
split horizon
IGRP 22-34
RIP 22-29
Stack Membership Discovery Protocol 10-14
Standby Command Configuration window 5-23

standby command switch
configuring 5-22
considerations 5-14
defined 5-2
priority 5-12
requirements 5-3
virtual IP address 5-13
See also cluster standby group and HSRP
standby group, cluster
See cluster standby group and HSRP
standby ip command 23-4
standby router 23-1
standby timers, HSRP 23-8
startup configuration
booting
manually 4-13
specific image 4-14
clearing B-19
configuration file
automatically downloading 4-12
specifying the filename 4-13
default boot configuration 4-12
static access mode 3-9
static access ports
assigning to VLAN 9-19
defined 8-2, 9-3
static addresses
See addresses
static IP routing 1-4
static MAC addressing 1-3
static routes, configuring 22-55
static routing 22-2
static VLAN membership 9-2
statistics
802.1X 7-14
CDP 13-5
interface 8-18
IP multicast routing 24-58
OSPF 22-44
statistics (continued)
QoS ingress and egress 20-56
RMON group Ethernet 16-5
RMON group history 16-5
SNMP input and output 18-10
VTP 9-13
storm control
configuring 12-3
definition 12-1
displaying 12-11
thresholds 12-1
STP
accelerating root port selection 10-12
BackboneFast, described 10-18
BPDU guard, described 10-11
BPDU message exchange 10-4
configuring
BackboneFast 10-36
BPDU guard 10-33
cross-stack UplinkFast 10-35
disable 10-22
forward-delay time 10-29
hello time 10-29
in cascaded stack 10-30
maximum aging time 10-30
path cost 10-27
Port Fast 10-32
port priority 10-26
root guard 10-36
root switch 10-22
secondary root switch 10-24
switch priority 10-28
UplinkFast with redundant links 10-34
cross-stack UplinkFast, described 10-13
default configuration 10-21
designated switch, defined 10-4
detecting indirect link failures 10-18
displaying status 10-31

STP (continued)

- EtherChannel guard
  - described 10-20
  - enabling 10-37
- extended system ID
  - affects on root switch 10-23
  - affects on the secondary root switch 10-24
  - overview 10-3
  - unexpected behavior 10-23
- features 1-3
- interface state, blocking to forwarding 10-10
- interface states
  - blocking 10-7
  - disabled 10-8
  - forwarding 10-7, 10-8
  - learning 10-7
  - listening 10-7
  - overview 10-6
- limitations with 802.1Q trunks 10-8
- load sharing
  - overview 9-29
  - using path costs 9-32
  - using port priorities 9-30
- loops 9-24
- multicast addresses, affect of 10-8
- overview 10-2
- path costs 9-32
- Port Fast
  - configuring 10-32
  - described 10-10
  - mode 9-36
- port priorities 9-30
- preventing root switch selection 10-20
- redundant connectivity 10-9
- root guard, described 1-3, 10-20
- root port, defined 10-4

STP (continued)

- root switch
  - affects of extended system ID 10-3, 10-23
  - configuring 10-23
  - election 10-3
  - unexpected behavior 10-23
- settings in a cascaded stack 10-30
- shutdown Port Fast-configured interface 10-11
- supported number of spanning-tree instances 10-2
- timers, described 10-5
- UplinkFast, described 10-12
- VLAN-bridge 10-9

stratum, NTP 6-32
stub areas, OSPF 22-39
subnet mask 22-5
subnet zero 22-8
summer time 6-44
SunNet Manager 1-6
supernet 22-8
SVIs
- and IP unicast routing 22-3
- and router ACLs 19-2
- connecting VLANs 8-5
- defined 8-4
- routing between VLANs 9-2

switch clustering technology 1-1, 1-6, 5-1
- See also CMS and clusters

switch clusters, planning considerations
- LRE profiles 5-17
- switch-specific features 5-18

switch console port 1-2
switched packets, ACLs on 19-39
switched ports 8-2
Switch Manager 3-2, 3-33
- See also Device Manager

switchport block multicast command 12-6
switchport block unicast command 12-6
switchport command 8-12
switchport protected command 12-5

switch priority, STP 10-28
switch software features 1-1
switch virtual interface
  See SVI
syslog
  See system message logging
system clock
  configuring
    daylight saving time 6-44
    manually 6-42
    summer time 6-44
    time zones 6-43
  displaying the time and date 6-42
  overview 6-32
  See also NTP
System Database Management
  See SDM
system message logging
  default configuration 17-3
  defining error message severity levels 17-8
  disabling 17-4
  displaying the configuration 17-12
  enabling 17-4
  facility keywords, described 17-12
  level keywords, described 17-9
  limiting messages 17-10
  message format 17-2
  overview 17-1
  sequence numbers, enabling and disabling 17-8
  setting the display destination device 17-4
  synchronizing log messages 17-6
  syslog facility 1-5
  timestamps, enabling and disabling 17-7
  UNIX syslog servers
    configuring the daemon 17-11
    configuring the logging facility 17-11
    facilities supported 17-12
system messages 3-19
system name
  default configuration 6-46
  default setting 6-46
  manual configuration 6-46
  See also DNS
system prompt
  default setting 6-46
  manual configuration 6-47
system resource templates 6-57
system routes, IGRP 22-30

T

tables, CMS 3-29
tabs, CMS 3-29
TAC
  inquiries xxxiii
  toll-free telephone numbers xxxiv
  website xxxiii
TACACS+
  accounting, defined 6-11
  authentication, defined 6-11
  authorization, defined 6-11
  configuring
    accounting 6-17
    authentication key 6-13
    authorization 6-16
    login authentication 6-14
  default configuration 6-13
  displaying the configuration 6-17
  identifying the server 6-13
  in clusters 5-17
  limiting the services to the user 6-16
  operation of 6-12
  overview 6-10
  support for 1-4
  tracking services accessed by user 6-17

tail drop
described 20-13
support for 1-4
tar files
creating B-6
displaying the contents of B-6
extracting B-7
image file format B-20
technical assistance
Cisco.com xxxiii
TAC xxxiii
TAC website xxxiii
toll-free telephone numbers xxxiv
Telnet
number of connections 1-2
setting a password 6-6
templates, system resources 6-57
Terminal Access Controller Access Control System Plus
See TACACS+
terminal lines, setting a password 6-6
TFTP
configuration files
downloading B-11
preparing the server B-10
uploading B-12
configuration files in base directory 4-6
configuring for autoconfiguration 4-5
image files
deleting B-23
downloading B-22
preparing the server B-22
uploading B-24
limiting access by servers 18-9
TFTP server 1-2
threshold, traffic level 12-2
time
See NTP and system clock
time-range command 19-15
time ranges in ACLs 19-15
timestamps in log messages 17-7
time zones 6-43
Token Ring VLANs
support for 9-15
TrCRF and TrBRF 9-6
toolbar 3-21
tool tips 3-27
Topology view
Collapse Cluster view 3-11
colors 3-14
described 3-10
device icons 3-12
device labels 3-13
display options 3-14
Expand Cluster view 3-11
icons 3-12
link icons 3-13
link labels 3-13
multilink icon 3-23
neighboring devices 3-12
pop-up menus 3-23
TOS 1-4
traceroute command 27-13
See also IP traceroute
traffic
blocking flooded 12-6
fragmented 19-4
unfragmented 19-4
traffic policing 1-4
traffic suppression 12-1
transparent mode, VTP 9-4, 9-11
trap-door mechanism 4-2
traps
configuring MAC address notification 6-54
configuring managers 18-7
defined 18-2
enabling 6-54, 18-7
notification types 18-7
overview 18-1, 18-3

troubleshooting
connectivity problems 27-11
detecting
EtherChannel misconfigurations 10-37
unidirectional links 14-1
determining packet disposition 27-15
displaying crash information 27-17
PIMv1 and PIMv2 interoperability problems 24-28
show forward command 27-15
with CiscoWorks 18-3
with debug commands 27-14
with ping 27-11
with system message logging 17-1
with traceroute 27-12
trunking encapsulation 1-3
trunk ports
configuring 9-25
defined 8-3, 9-3
encapsulation 9-25, 9-31, 9-32
trunks
allowed-VLAN list 9-27
configuring 9-25, 9-31, 9-32
encapsulation 9-24
Fast Ethernet 9-22
Gigabit Ethernet 9-22
interacting with other features 9-24
ISL 9-22
load sharing
setting STP path costs 9-32
using STP port priorities 9-30
native VLAN for untagged traffic 9-29
parallel 9-32
pruning-eligible list 9-28
to non-DTP device 9-23
understanding 9-22
twisted-pair Ethernet, detecting unidirectional links 14-1
type-of-service
See TOS

U

UDLD
default configuration 14-3
echoing detection mechanism 14-2
enabling
globally 14-3
per interface 14-4
link-detection mechanism 14-1
neighbor database 14-2
overview 14-1
resetting an interface 14-4
status, displaying 14-5
support for 1-3
UDP, configuring 22-19
unauthorized ports with 802.1X 7-4
unequal-cost load balancing, IGRP 22-31
unicast storm control command 12-3
unicast storms 12-1
unicast traffic, blocking 12-6
UniDirectional Link Detection protocol
See UDLD
UNIX syslog servers
daemon configuration 17-11
facilities supported 17-12
message logging configuration 17-11
unrecognized Type-Length-Value (TLV) support 9-6
upgrading software images
See downloading
UplinkFast
described 10-12
redundant links 10-34
support for 1-3

uploading
configuration files
preparing B-10, B-13, B-16
reasons for B-8
using FTP B-15
using RCP B-18
using TFTP B-12
image files
preparing B-22, B-25, B-29
reasons for B-19
using FTP B-28
using RCP B-32
using TFTP B-24
URLs, Cisco xxxi
User Datagram Protocol
See UDP
user EXEC mode 2-2
username-based authentication 6-7

V

verifying changes in CMS 3-32
version-dependent transparent mode 9-6
virtual IP address
cluster standby group 5-13, 5-22
command switch 5-13, 5-22
See also IP addresses
virtual router 23-1, 23-2
VLAN ACLs
See VLAN maps
VLAN configuration mode 2-2
VLAN management domain 9-4
VLAN Management Policy Server
See VMPS
VLAN map entries, order of 19-28
VLAN maps
applying 19-32
common uses for 19-33
configuration example 19-34
configuration guidelines 19-28
configuring 19-27
creating 19-30
defined 19-2
denying access example 19-35
denying and permitting packets 19-30
displaying 19-33
examples 19-35
support for 1-3
usage 19-3
with router ACLs 19-42
VLAN membership
confirming 9-39
modes 3-9, 9-3
VLAN port membership modes 3-9
VLAN Query Protocol (VQP) 9-33
VLANs
adding to VTP database 9-17
aging dynamic addresses 10-10
allowed on trunk 9-27
changing 9-18
configuration guidelines 9-16
configuring 9-1
configuring in VTP database 9-17
connecting through SVIs 8-5
default configuration 9-15
deleting from VTP database 9-18
described 8-2, 9-1
features 1-3
illustrated 9-2
in the VTP database 9-17
ISL 9-22
limiting source traffic with SPAN 15-12
modifying 9-18
monitoring with SPAN 15-11

VLANs (continued)

native, configuring 9-29
number supported 1-3, 9-2
port membership modes 9-3
static-access ports 9-18, 9-19
STP and 802.1Q trunks 10-8
supported 9-2
Token Ring 9-15
VLAN-bridge STP 10-9, 26-1
VTP database and 9-15
VTP modes 9-4

VLAN Trunking Protocol
See VTP

VLAN trunks 9-22

VMPS
administering 9-40
configuration guidelines 9-36
default configuration 9-37
dynamic port membership
example 9-40
overview 9-34
reconfirming 9-39
troubleshooting 9-40
entering server address 9-37
mapping MAC addresses to VLANs 9-33
monitoring 9-40
reconfirmation interval, changing 9-39
reconfirming membership 9-39
retry count, changing 9-39

voice VLANs
See VVIDs

VQP 1-3

VTP
advertisements 9-5, 9-25
client mode, configuring 9-11
configuration
guidelines 9-8
requirements 9-9

VTP (continued)

configuration requirements 9-9
configuring
client mode 9-11
server mode 9-10
transparent mode 9-11
consistency checks 9-6
database 9-17
configuring VLANs 9-17
displaying 9-21
parameters 9-15
default configuration 9-8
described 9-3
disabling 9-11
domain names 9-8
domains 9-4
modes
client 9-4, 9-11
server 9-4, 9-10
transitions 9-4
transparent 9-4, 9-11
monitoring 9-13
passwords 9-9
pruning
disabling 9-13
enabling 9-13
examples 9-7
overview 9-6
pruning-eligible list, changing 9-28
server mode, configuring 9-10
statistics 9-13
support for 1-3
Token Ring support 9-6
transparent mode, configuring 9-11
using 9-3
version, guidelines 9-9
version 1 9-6

VTP (continued)
version 2
configuration guidelines 9-9
disabling 9-12
enabling 9-12
overview 9-6
VLAN parameters 9-15
VTP monitoring 9-13
VTP pruning 1-3
VVIDs 1-8

W

web-based management software
See CMS
Weighted Random Early Detection
See WRED
Weighted Round Robin
See WRR
weighted round robin, described 20-3
window components, CMS 3-28
wizards 1-6, 3-26
WRED 1-4, 20-14
WRR 1-4, 20-3

X

XMODEM protocol 27-2